

Internet Group Management Protocol, Version 2

Status of this Memo

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

Copyright Notice

Copyright (C) The Internet Society (1997). All Rights Reserved.

Abstract

This memo documents IGMPv2, used by IP hosts to report their multicast group memberships to routers. It updates STD 5, RFC 1112.

IGMPv2 allows group membership termination to be quickly reported to the routing protocol, which is important for high-bandwidth multicast groups and/or subnets with highly volatile group membership.

This document is a product of the Inter-Domain Multicast Routing working group within the Internet Engineering Task Force. Comments are solicited and should be addressed to the working group's mailing list at idmr@cs.ucl.ac.uk and/or the author(s).

1. Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC 2119].

2. Introduction

The Internet Group Management Protocol (IGMP) is used by IP hosts to report their multicast group memberships to any immediately-neighborings multicast routers. This memo describes only the use of IGMP between hosts and routers to determine group membership. Routers that are members of multicast groups are expected to behave

as hosts as well as routers, and may even respond to their own queries. IGMP may also be used between routers, but such use is not specified here.

Like ICMP, IGMP is an integral part of IP. It is required to be implemented by all hosts wishing to receive IP multicasts. IGMP messages are encapsulated in IP datagrams, with an IP protocol number of 2. All IGMP messages described in this document are sent with IP TTL 1, and contain the IP Router Alert option [RFC 2113] in their IP header. All IGMP messages of concern to hosts have the following format:

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Type          | Max Resp Time |           Checksum           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Group Address                          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

2.1. Type

There are three types of IGMP messages of concern to the host-router interaction:

0x11 = Membership Query

There are two sub-types of Membership Query messages:

- General Query, used to learn which groups have members on an attached network.
- Group-Specific Query, used to learn if a particular group has any members on an attached network.

These two messages are differentiated by the Group Address, as described in section 1.4. Membership Query messages are referred to simply as "Query" messages.

0x16 = Version 2 Membership Report

0x17 = Leave Group

There is an additional type of message, for backwards-compatibility with IGMPv1:

0x12 = Version 1 Membership Report

This document refers to Membership Reports simply as "Reports". When no version is specified, the statement applies equally to both versions.

Unrecognized message types should be silently ignored. New message types may be used by newer versions of IGMP, by multicast routing protocols, or other uses.

2.2. Max Response Time

The Max Response Time field is meaningful only in Membership Query messages, and specifies the maximum allowed time before sending a responding report in units of 1/10 second. In all other messages, it is set to zero by the sender and ignored by receivers.

Varying this setting allows IGMPv2 routers to tune the "leave latency" (the time between the moment the last host leaves a group and when the routing protocol is notified that there are no more members), as discussed in section 7.8. It also allows tuning of the burstiness of IGMP traffic on a subnet, as discussed in section 7.3.

2.3. Checksum

The checksum is the 16-bit one's complement of the one's complement sum of the whole IGMP message (the entire IP payload). For computing the checksum, the checksum field is set to zero. When transmitting packets, the checksum MUST be computed and inserted into this field. When receiving packets, the checksum MUST be verified before processing a packet.

2.4. Group Address

In a Membership Query message, the group address field is set to zero when sending a General Query, and set to the group address being queried when sending a Group-Specific Query.

In a Membership Report or Leave Group message, the group address field holds the IP multicast group address of the group being reported or left.

2.5. Other fields

Note that IGMP messages may be longer than 8 octets, especially future backwards-compatible versions of IGMP. As long as the Type is one that is recognized, an IGMPv2 implementation MUST ignore anything past the first 8 octets while processing the packet. However, the IGMP checksum is always computed over the whole IP payload, not just over the first 8 octets.

3. Protocol Description

Note that defaults for timer values are described later in this document. Timer and counter names appear in square brackets.

The term "interface" is sometimes used in this document to mean "the primary interface on an attached network"; if a router has multiple physical interfaces on a single network this protocol need only run on one of them. Hosts, on the other hand, need to perform their actions on all interfaces that have memberships associated with them.

Multicast routers use IGMP to learn which groups have members on each of their attached physical networks. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership. "Multicast group memberships" means the presence of at least one member of a multicast group on a given attached network, not a list of all of the members. With respect to each of its attached networks, a multicast router may assume one of two roles: Querier or Non-Querier. There is normally only one Querier per physical network. All multicast routers start up as a Querier on each attached network. If a multicast router hears a Query message from a router with a lower IP address, it MUST become a Non-Querier on that network. If a router has not heard a Query message from another router for [Other Querier Present Interval], it resumes the role of Querier. Routers periodically [Query Interval] send a General Query on each attached network for which this router is the Querier, to solicit membership information. On startup, a router SHOULD send [Startup Query Count] General Queries spaced closely together [Startup Query Interval] in order to quickly and reliably determine membership information. A General Query is addressed to the all-systems multicast group (224.0.0.1), has a Group Address field of 0, and has a Max Response Time of [Query Response Interval].

When a host receives a General Query, it sets delay timers for each group (excluding the all-systems group) of which it is a member on the interface from which it received the query. Each timer is set to a different random value, using the highest clock granularity available on the host, selected from the range (0, Max Response Time] with Max Response Time as specified in the Query packet. When a host receives a Group-Specific Query, it sets a delay timer to a random value selected from the range (0, Max Response Time] as above for the group being queried if it is a member on the interface from which it received the query. If a timer for the group is already running, it is reset to the random value only if the requested Max Response Time is less than the remaining value of the running timer. When a group's timer expires, the host multicasts a Version 2 Membership Report to the group, with IP TTL of 1. If the host receives another

host's Report (version 1 or 2) while it has a timer running, it stops its timer for the specified group and does not send a Report, in order to suppress duplicate Reports.

When a router receives a Report, it adds the group being reported to the list of multicast group memberships on the network on which it received the Report and sets the timer for the membership to the [Group Membership Interval]. Repeated Reports refresh the timer. If no Reports are received for a particular group before this timer has expired, the router assumes that the group has no local members and that it need not forward remotely-originated multicasts for that group onto the attached network.

When a host joins a multicast group, it should immediately transmit an unsolicited Version 2 Membership Report for that group, in case it is the first member of that group on the network. To cover the possibility of the initial Membership Report being lost or damaged, it is recommended that it be repeated once or twice after short delays [Unsolicited Report Interval]. (A simple way to accomplish this is to send the initial Version 2 Membership Report and then act as if a Group-Specific Query was received for that group, and set a timer appropriately).

When a host leaves a multicast group, if it was the last host to reply to a Query with a Membership Report for that group, it SHOULD send a Leave Group message to the all-routers multicast group (224.0.0.2). If it was not the last host to reply to a Query, it MAY send nothing as there must be another member on the subnet. This is an optimization to reduce traffic; a host without sufficient storage to remember whether or not it was the last host to reply MAY always send a Leave Group message when it leaves a group. Routers SHOULD accept a Leave Group message addressed to the group being left, in order to accommodate implementations of an earlier version of this standard. Leave Group messages are addressed to the all-routers group because other group members have no need to know that a host has left the group, but it does no harm to address the message to the group.

When a Querier receives a Leave Group message for a group that has group members on the reception interface, it sends [Last Member Query Count] Group-Specific Queries every [Last Member Query Interval] to the group being left. These Group-Specific Queries have their Max Response time set to [Last Member Query Interval]. If no Reports are received after the response time of the last query expires, the routers assume that the group has no local members, as above. Any Querier to non-Querier transition is ignored during this time; the same router keeps sending the Group-Specific Queries.

Non-Queriers MUST ignore Leave Group messages, and Queriers SHOULD ignore Leave Group messages for which there are no group members on the reception interface.

When a non-Querier receives a Group-Specific Query message, if its existing group membership timer is greater than [Last Member Query Count] times the Max Response Time specified in the message, it sets its group membership timer to that value.

4. Compatibility with IGMPv1 Routers

An IGMPv2 host may be placed on a subnet where the Querier router has not yet been upgraded to IGMPv2. The following requirements apply:

The IGMPv1 router will send General Queries with the Max Response Time set to 0. This MUST be interpreted as a value of 100 (10 seconds).

The IGMPv1 router expects Version 1 Membership Reports in response to its Queries, and will not pay attention to Version 2 Membership Reports. Therefore, a state variable MUST be kept for each interface, describing whether the multicast Querier on that interface is running IGMPv1 or IGMPv2. This variable MUST be based upon whether or not an IGMPv1 query was heard in the last [Version 1 Router Present Timeout] seconds, and MUST NOT be based upon the type of the last Query heard. This state variable MUST be used to decide what type of Membership Reports to send for unsolicited Membership Reports as well as Membership Reports in response to Queries.

An IGMPv2 host MAY suppress Leave Group messages on a network where the Querier is using IGMPv1.

An IGMPv2 router may be placed on a subnet where at least one router on the subnet has not yet been upgraded to IGMPv2. The following requirements apply:

If any IGMPv1 routers are present, the querier MUST use IGMPv1. The use of IGMPv1 must be administratively configured, as there is no reliable way of dynamically determining whether IGMPv1 routers are present on a network. Implementations MAY provide a way for system administrators to enable the use of IGMPv1 on their routers; in the absence of explicit configuration, the configuration MUST default to IGMPv2. When in IGMPv1 mode, routers MUST send Periodic Queries with a Max Response Time of 0, and MUST ignore Leave Group messages. They SHOULD also warn about receiving an IGMPv2 query, although such warnings MUST be rate-limited.

If a router is not explicitly configured to use IGMPv1 and hears an IGMPv1 Query, it SHOULD log a warning. These warnings MUST be rate-limited.

5. Compatibility with IGMPv1 Hosts

An IGMPv2 host may be placed on a subnet where there are hosts that have not yet been upgraded to IGMPv2. The following requirement applies:

The host MUST allow its Membership Report to be suppressed by either a Version 1 Membership Report or a Version 2 Membership Report.

An IGMPv2 router may be placed on a subnet where there are hosts that have not yet been upgraded to IGMPv2. The following requirements apply:

If a router receives a Version 1 Membership Report, it MUST set a timer to note that there are version 1 hosts present which are members of the group for which it heard the report. This timer should be the same as the [Group Membership Interval].

If there are version 1 hosts present for a particular group, a router MUST ignore any Leave Group messages that it receives for that group.

6. Host State Diagram

Host behavior is more formally specified by the state transition diagram below. A host may be in one of three possible states with respect to any single IP multicast group on any single network interface:

- "Non-Member" state, when the host does not belong to the group on the interface. This is the initial state for all memberships on all network interfaces; it requires no storage in the host.
- "Delaying Member" state, when the host belongs to the group on the interface and has a report delay timer running for that membership.
- "Idle Member" state, when the host belongs to the group on the interface and does not have a report delay timer running for that membership.

There are five significant events that can cause IGMP state transitions:

- "join group" occurs when the host decides to join the group on the interface. It may occur only in the Non-Member state.
- "leave group" occurs when the host decides to leave the group on the interface. It may occur only in the Delaying Member and Idle Member states.
- "query received" occurs when the host receives either a valid General Membership Query message, or a valid Group-Specific Membership Query message. To be valid, the Query message must be at least 8 octets long, and have a correct IGMP checksum. The group address in the IGMP header must either be zero (a General Query) or a valid multicast group address (a Group-Specific Query). A General Query applies to all memberships on the interface from which the Query is received. A Group-Specific Query applies to membership in a single group on the interface from which the Query is received. Queries are ignored for memberships in the Non-Member state.
- "report received" occurs when the host receives a valid IGMP Membership Report message (Version 1 or Version 2). To be valid, the Report message must be at least 8 octets long and have a correct IGMP checksum. A Membership Report applies only to the membership in the group identified by the Membership Report, on the interface from which the Membership Report is received. It is ignored for memberships in the Non-Member or Idle Member state.
- "timer expired" occurs when the report delay timer for the group on the interface expires. It may occur only in the Delaying Member state.

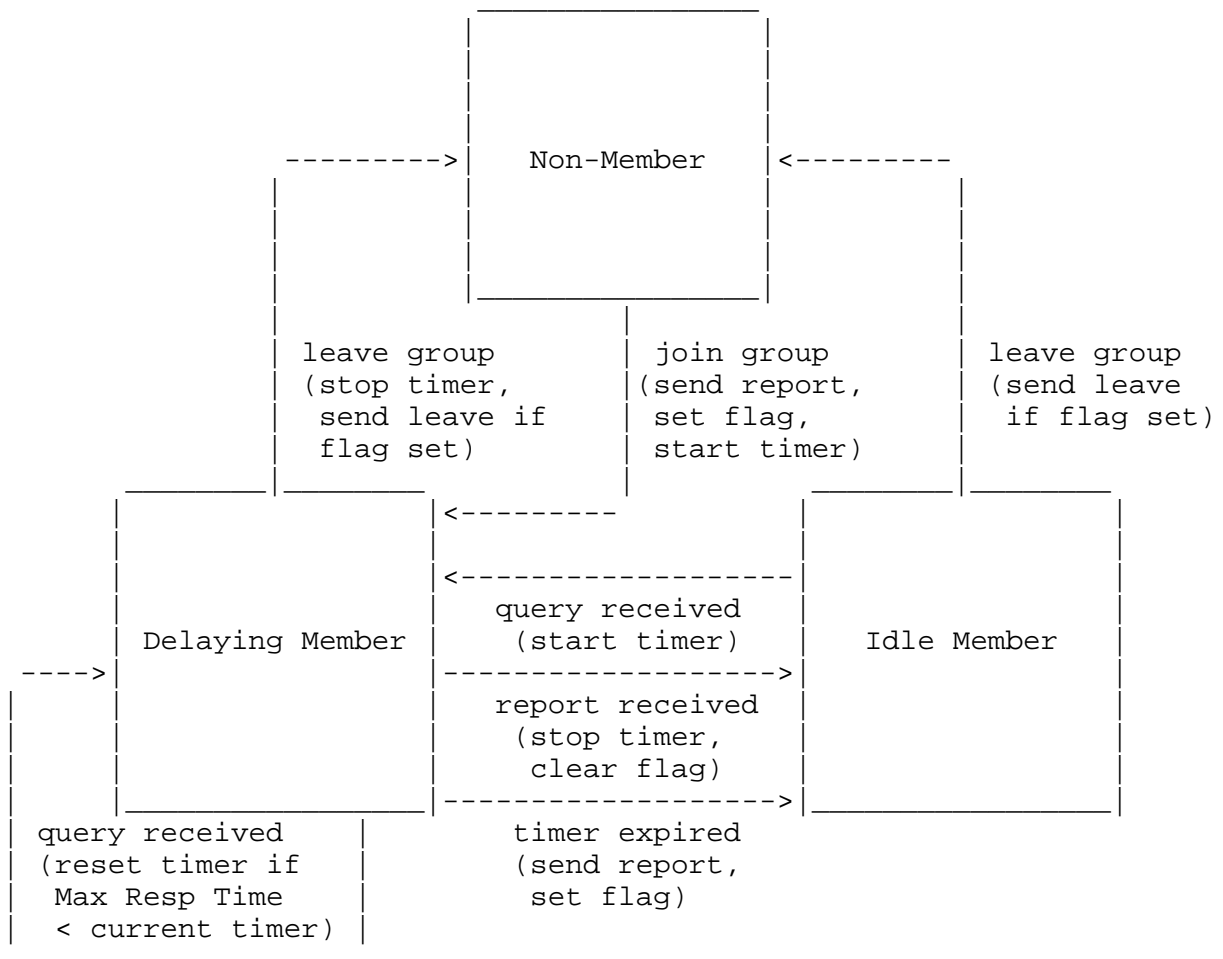
All other events, such as receiving invalid IGMP messages, or IGMP messages other than Query or Report, are ignored in all states.

There are seven possible actions that may be taken in response to the above events:

- "send report" for the group on the interface. The type of report is determined by the state of the interface. The Report Message is sent to the group being reported.

- "send leave" for the group on the interface. If the interface state says the Querier is running IGMPv1, this action SHOULD be skipped. If the flag saying we were the last host to report is cleared, this action MAY be skipped. The Leave Message is sent to the ALL-ROUTERS group (224.0.0.2).
- "set flag" that we were the last host to send a report for this group.
- "clear flag" since we were not the last host to send a report for this group.
- "start timer" for the group on the interface, using a delay value chosen uniformly from the interval (0, Max Response Time], where Max Response time is specified in the Query. If this is an unsolicited Report, the timer is set to a delay value chosen uniformly from the interval (0, [Unsolicited Report Interval]).
- "reset timer" for the group on the interface to a new value, using a delay value chosen uniformly from the interval (0, Max Response Time], as described in "start timer".
- "stop timer" for the group on the interface.

In all of the following state diagrams, each state transition arc is labeled with the event that causes the transition, and, in parentheses, any actions taken during the transition. Note that the transition is always triggered by the event; even if the action is conditional, the transition still occurs.



The all-systems group (address 224.0.0.1) is handled as a special case. The host starts in Idle Member state for that group on every interface, never transitions to another state, and never sends a report for that group.

In addition, a host may be in one of two possible states with respect to any single network interface:

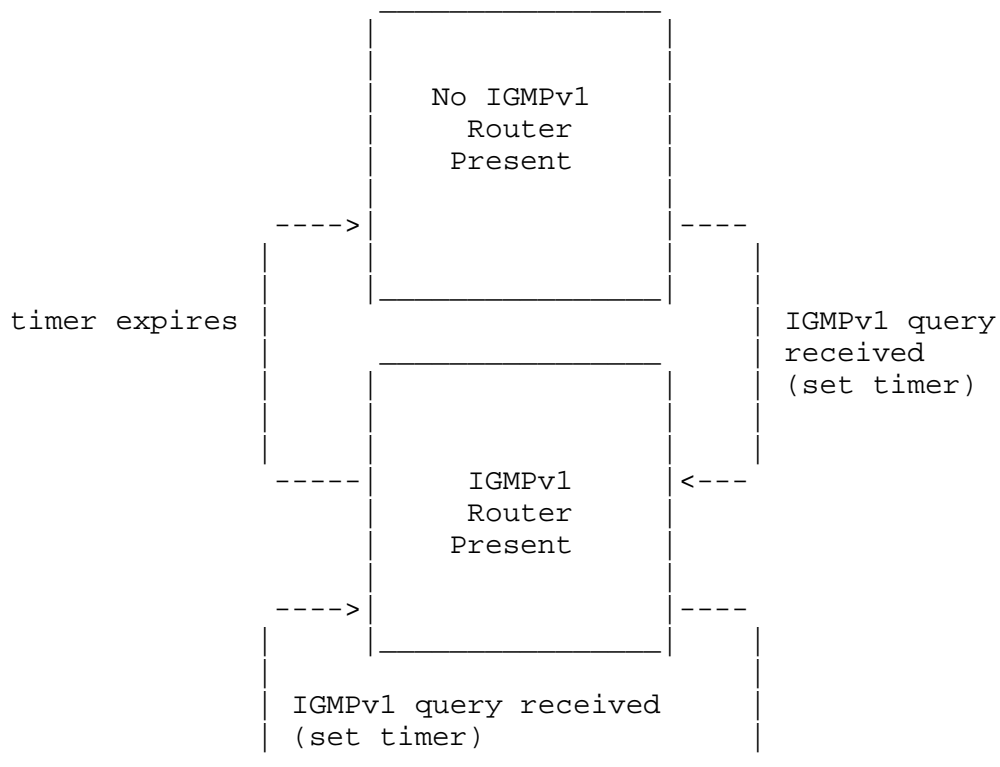
- "No IGMPv1 Router Present", when the host has not heard an IGMPv1 style query for the [Version 1 Router Present Timeout]. This is the initial state.
- "IGMPv1 Router Present", when the host has heard an IGMPv1 style query within the [Version 1 Router Present Timeout].

There are two events that can cause state transitions:

- "IGMPv1 query received", when the host receives a query with the Max Response Time field set to 0.
- "timer expires", when the timer set to note the presence of an IGMPv1 router expires.

And a single action that can be triggered by an event:

- "set timer", setting the timer to its maximum value [Version 1 Router Present Timeout] and (re)starting it.



7. Router State Diagram

Router behavior is more formally specified by the state transition diagrams below.

A router may be in one of two possible states with respect to any single attached network:

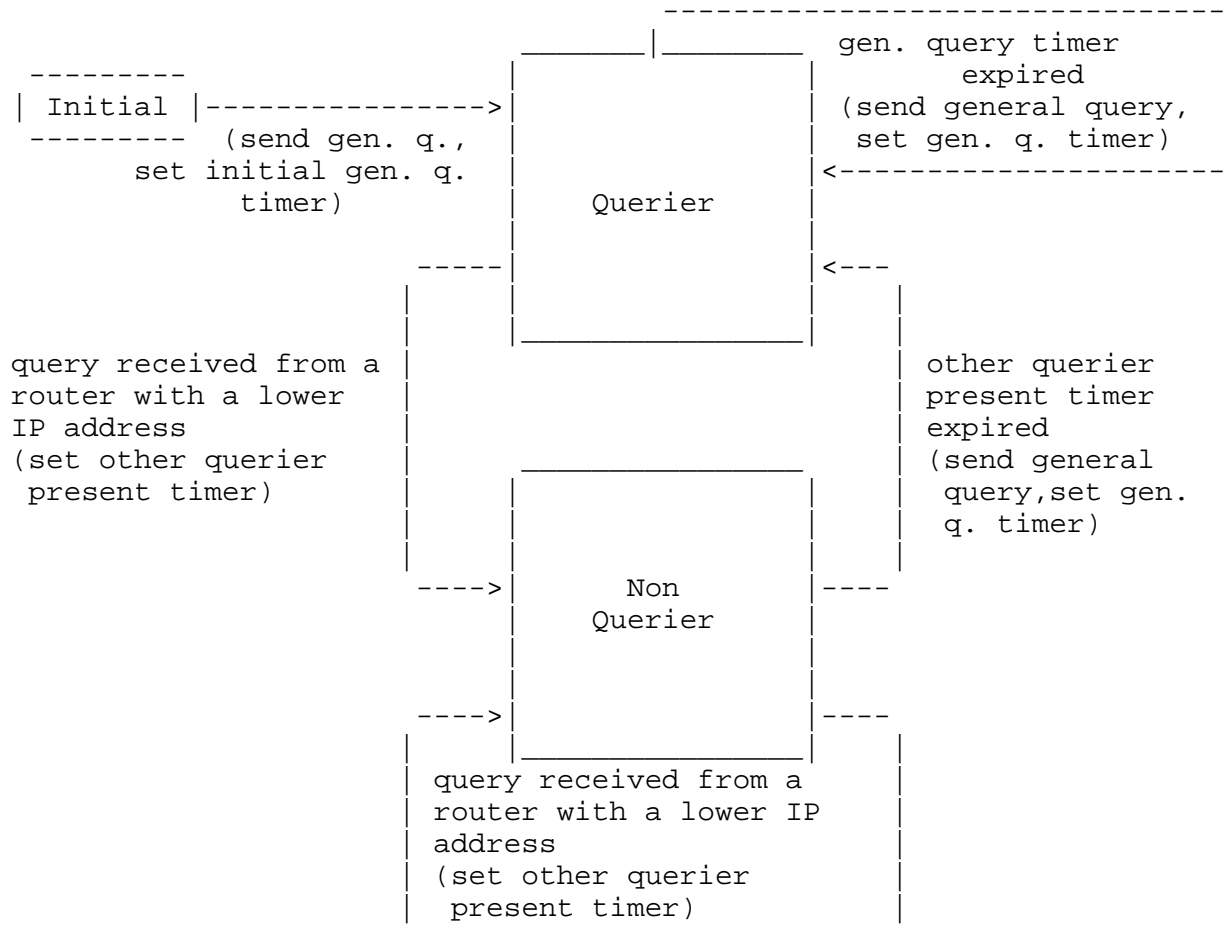
- "Querier", when this router is designated to transmit IGMP Membership Queries on this network.
- "Non-Querier", when there is another router designated to transmit IGMP membership Queries on this network.

The following three events can cause the router to change states:

- "query timer expired" occurs when the timer set for query transmission expires.
- "query received from a router with a lower IP address" occurs when an IGMP Membership Query is received from a router on the same network with a lower IP address.
- "other querier present timer expired" occurs when the timer set to note the presence of another querier with a lower IP address on the network expires.

There are three actions that may be taken in response to the above events:

- "start general query timer" for the attached network.
- "start other querier present timer" for the attached network [Other Querier Present Interval].
- "send general query" on the attached network. The General Query is sent to the all-systems group (224.0.0.1), and has a Max Response Time of [Query Response Interval].



A router should start in the Initial state on all attached networks, and immediately move to Querier state.

In addition, to keep track of which groups have members, a router may be in one of four possible states with respect to any single IP multicast group on any single attached network:

- "No Members Present" state, when there are no hosts on the network which have sent reports for this multicast group. This is the initial state for all groups on the router; it requires no storage in the router.
- "Members Present" state, when there is a host on the network which has sent a Membership Report for this multicast group.

- "Version 1 Members Present" state, when there is an IGMPv1 host on the network which has sent a Version 1 Membership Report for this multicast group.
- "Checking Membership" state, when the router has received a Leave Group message but has not yet heard a Membership Report for the multicast group.

There are six significant events that can cause router state transitions:

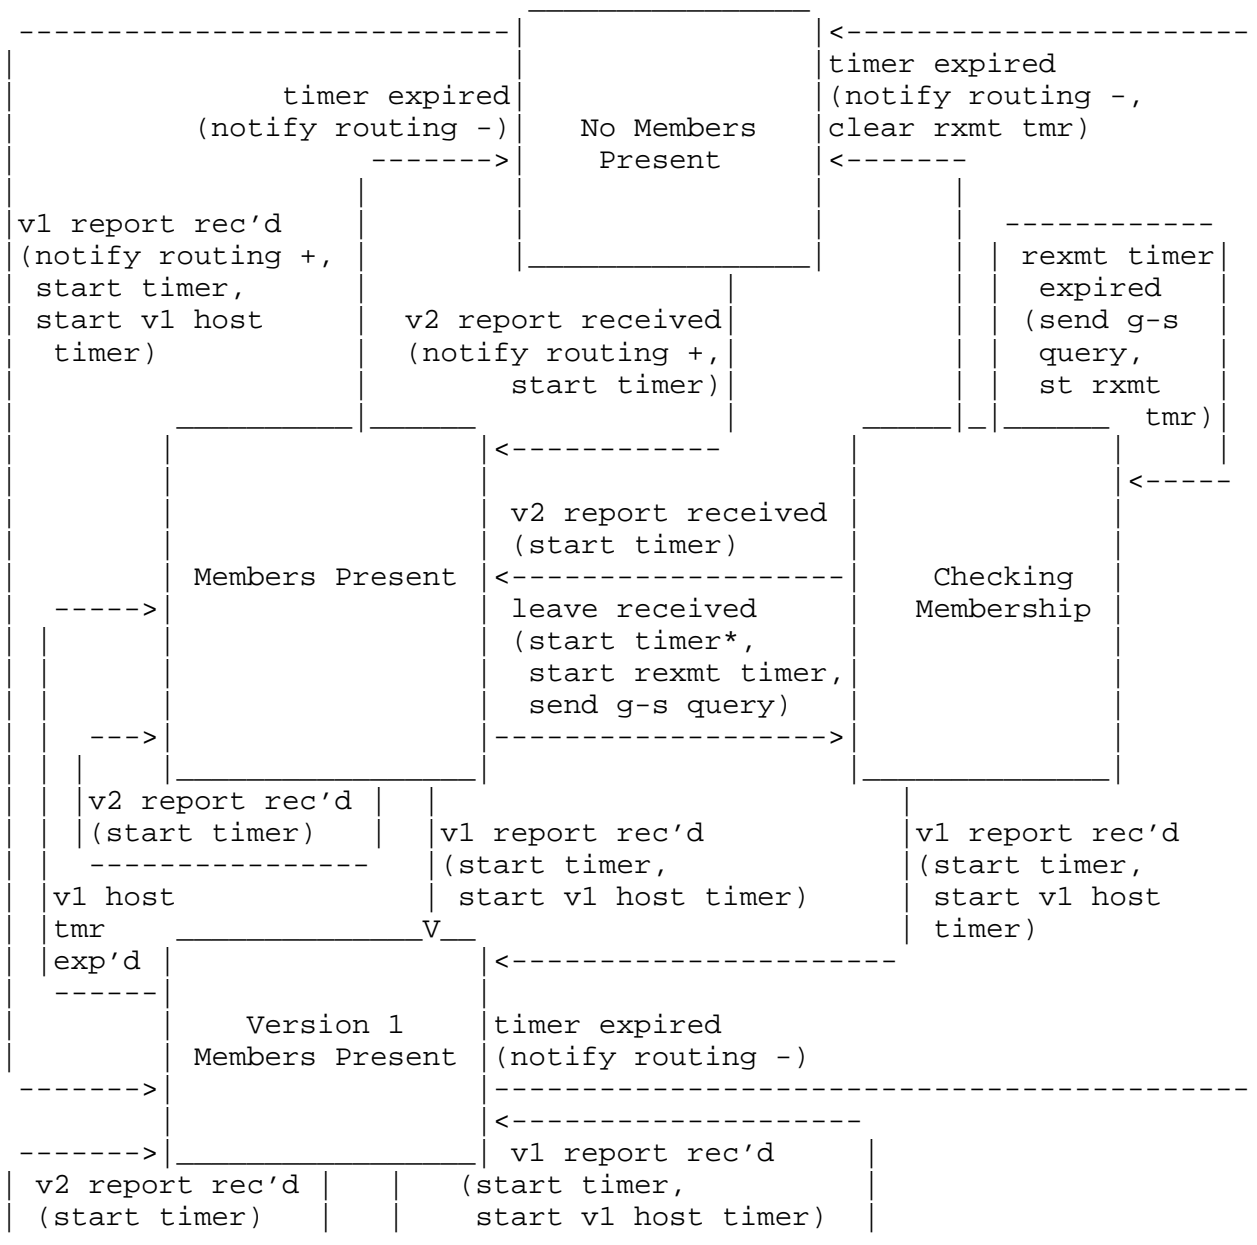
- "v2 report received" occurs when the router receives a Version 2 Membership Report for the group on the interface. To be valid, the Report message must be at least 8 octets long and must have a correct IGMP checksum.
- "v1 report received" occurs when the router receives a Version 1 Membership report for the group on the interface. The same validity requirements apply.
- "leave received" occurs when the router receives an IGMP Group Leave message for the group on the interface. To be valid, the Leave message must be at least 8 octets long and must have a correct IGMP checksum.
- "timer expired" occurs when the timer set for a group membership expires.
- "retransmit timer expired" occurs when the timer set to retransmit a group-specific Membership Query expires.
- "v1 host timer expired" occurs when the timer set to note the presence of version 1 hosts as group members expires.

There are six possible actions that may be taken in response to the above events:

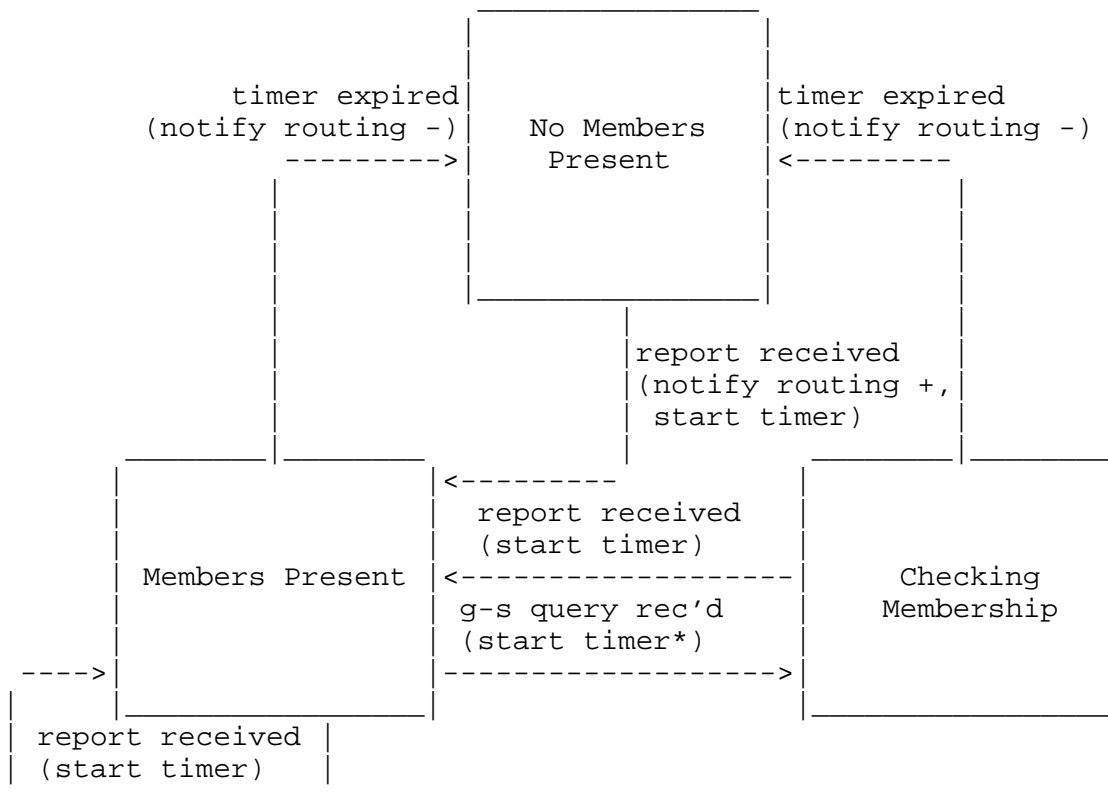
- "start timer" for the group membership on the interface - also resets the timer to its initial value [Group Membership Interval] if the timer is currently running.
- "start timer*" for the group membership on the interface - this alternate action sets the timer to [Last Member Query Interval] * [Last Member Query Count] if this router is a Querier, or the [Max Response Time] in the packet * [Last Member Query Count] if this router is a non-Querier.

- "start retransmit timer" for the group membership on the interface [Last Member Query Interval].
- "start vl host timer" for the group membership on the interface, also resets the timer to its initial value [Group Membership Interval] if the timer is currently running.
- "send group-specific query" for the group on the attached network. The Group-Specific Query is sent to the group being queried, and has a Max Response Time of [Last Member Query Interval].
- "notify routing +" notify the routing protocol that there are members of this group on this connected network.
- "notify routing -" notify the routing protocol that there are no longer any members of this group on this connected network.

The state diagram for a router in Querier state follows:



The state diagram for a router in Non-Querier state is similar, but non-Queriers do not send any messages and are only driven by message reception. Note that non-Queriers do not care whether a Membership Report message is Version 1 or Version 2.



8. List of timers and default values

Most of these timers are configurable. If non-default settings are used, they MUST be consistent among all routers 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 subnet. If a subnet is expected to be lossy, the Robustness Variable may be increased. IGMP 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 IGMP messages on the subnet; larger values cause IGMP Queries to be sent less often.

8.3. Query Response Interval

The Max Response Time inserted into the periodic General Queries. Default: 100 (10 seconds)

By varying the [Query Response Interval], an administrator may tune the burstiness of IGMP messages on the subnet; 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. Group Membership Interval

The Group Membership Interval is the amount of time that must pass before a multicast router decides there are no more members of a group on a network. This value MUST be ((the Robustness Variable) times (the Query Interval)) plus (one Query Response Interval).

8.5. Other Querier Present Interval

The Other Querier Present Interval 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 Member Query Interval

The Last Member Query Interval is the Max Response Time inserted into Group-Specific Queries sent in response to Leave Group messages, and is also the amount of time between Group-Specific Query messages.
Default: 10 (1 second)

This value may be tuned to modify the "leave latency" of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

8.9. Last Member Query Count

The Last Member Query Count is the number of Group-Specific Queries sent before the router assumes there are no local members. Default: the Robustness Variable.

8.10. Unsolicited Report Interval

The Unsolicited Report Interval is the time between repetitions of a host's initial report of membership in a group. Default: 10 seconds.

8.11. Version 1 Router Present Timeout

The Version 1 Router Present Timeout is how long a host must wait after hearing a Version 1 Query before it may send any IGMPv2 messages. Value: 400 seconds.

9. Message destinations

This information is provided elsewhere in the document, but is summarized here for convenience.

Message Type	Destination Group
-----	-----
General Query	ALL-SYSTEMS (224.0.0.1)
Group-Specific Query	The group being queried
Membership Report	The group being reported
Leave Message	ALL-ROUTERS (224.0.0.2)

Note: in older (i.e., non-standard and now obsolete) versions of IGMPv2, hosts send Leave Messages to the group being left. A router SHOULD accept Leave Messages addressed to the group being left in the interests of backwards compatibility with such hosts. In all cases, however, hosts MUST send to the ALL-ROUTERS address to be compliant with this specification.

10. Security Considerations

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

Query Message:

A forged Query message from a machine with a lower IP 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 Leave Messages, traffic might flow to groups with no members for up to [Group Membership Interval].

A forged Query message sent to a group with members will cause the hosts which are members of the group to report their memberships. This causes a small amount of extra traffic on the LAN, but causes no protocol problems.

Report messages:

A forged Report message may cause multicast routers to think there are members of a group on a subnet when there are not. Forged Report messages from the local subnet are meaningless, since joining a group 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:

- Ignore the Report if you cannot identify the source address of the packet as belonging to a subnet assigned to the interface on which the packet was received. This solution means that Reports sent by mobile hosts without addresses on the local subnet will be ignored.
- Ignore Report messages without Router Alert options [RFC 2113], 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). This solution breaks backwards compatibility with implementations of earlier versions of this specification which did not require Router Alert.

A forged Version 1 Report Message may put a router into "version 1 members present" state for a particular group, meaning that the router will ignore Leave messages. This can cause traffic to flow to groups with no members for up to [Group Membership Interval]. There are two defenses against forged v1 Reports:

- To defend against externally sourced v1 Reports, ignore the Report if you cannot identify the source address of the packet as belonging to a subnet assigned to the interface on which the packet was received. This solution means that v1 Reports sent by mobile hosts without addresses on the local subnet will be ignored.
- Provide 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 "fast leave" is critical. This solution protects against forged version 1 reports from the local subnet as well.

Leave message:

A forged Leave message will cause the Querier to send out Group-Specific Queries for the group in question. This causes extra processing on each router and on each member of the group, but can not cause loss of desired traffic. There are two defenses against externally forged Leave messages:

- Ignore the Leave message if you cannot identify the source address of the packet as belonging to a subnet assigned to the interface on which the packet was received. This solution means that Leave messages sent by mobile hosts without addresses on the local subnet will be ignored.
- Ignore Leave messages without Router Alert options [RFC 2113], and require that routers not forward Leave messages. (The requirement is not a requirement of generalized filtering in the forwarding path, since the packets already have Router Alert options in them). This solution breaks backwards compatibility with implementations of earlier versions of this specification which did not require Router Alert.

11. Acknowledgments

IGMPv2 was designed by Rosen Sharma and Steve Deering.

12. References

- RFC 2119 Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- RFC 2113 Katz, D., "IP Router Alert Option," RFC 2113, February 1997.
- RFC 1112 Deering, S., "Host Extensions for IP Multicasting", STD 5, RFC 1112, August 1989.

13. Appendix I - Changes from IGMPv1

The IGMPv1 "Version" and "Type" fields are combined into a single "Type" field.

A new IGMP Type is assigned to Version 2 Membership Report messages, so a router may tell the difference between an IGMPv1 and IGMPv2 host report.

A new IGMP Type is created for the IGMPv2 Leave Group message.

The Membership Query message is changed so that a previously unused field contains a new value, the Max Response Time.

The IGMPv2 spec now specifies a querier election mechanism. In IGMPv1, the querier election was left up to the multicast routing protocol, and different protocols used different mechanisms. This could result in more than one querier per network, so the election mechanism has been standardized in IGMPv2. However, this means that care must be taken when an IGMPv2 router is trying to coexist with an IGMPv1 router that uses a different querier election mechanism. In particular, it means that an IGMPv2 router must be able to act as an IGMPv1 router on a particular network if configured to do so. The actions required include:

- Set the Max Response Time field to 0 in all queries.
- Ignore Leave Group messages.

The IGMPv2 spec relaxes the requirements on validity-checking for Membership Queries and Membership Reports. When upgrading an implementation, be sure to remove any checks that do not belong.

The IGMPv2 spec requires the presence of the IP Router Alert option [RFC 2113] in all packets described in this memo.

14. Author's Address

William C. Fenner
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304
Phone: +1 650 812 4816

EMail: fenner@parc.xerox.com

15. Full Copyright Statement

Copyright (C) The Internet Society (1997). All Rights Reserved.

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

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

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

