ETSI TS 102 293 V1.1.1 (2004-02)

Technical Specification

Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM) services and architectures; IP Interworking over satellite; Multicast group management; IGMP adaptation

Reference

DTS/SES-00088

Keywords

broadband, interworking, IP, multimedia, satellite

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

Individual copies of the present document can be downloaded from:
http://www.etsi.org

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
http://portal.etsi.org/tb/status/status.asp

If you find errors in the present document, send your comment to:
editor@etsi.org

Copyright Notification

Contents

Intellectual Property Rights 5

Foreword 5

Introduction 5

1 Scope 6

2 References 6

3 Definitions and abbreviations 6
3.1 Definitions 6
3.2 Abbreviations 7

4 Overview and summary 7
4.1 Background 7
4.2 Need for IGMP adaptation 7
4.3 IGMP modes 8
4.4 Advantages of S-IGMP solution 8
4.5 Influence of IGMP proxies, etc. 8

5 Applicability and assumptions 8
5.1 Use of IGMP in BSM multicast architectures 8
5.2 IGMP modes and multicast routing protocol options 9
5.3 IGMP network architecture applicable to BSM 10
5.4 IGMP-specific BSM functional model 10
5.5 IGMP issues in a satellite system 11
5.6 IGMP messages 11

6 S-IGMP specification 12
6.1 General 12
6.2 S-IGMP client 12
6.3 S-IGMP router (querier) 13
6.3.1 Overall querier state diagram 13
6.3.2 Router state diagram in querier mode 14
6.3.2.1 Parameter definitions 15
6.3.3 Detailed S-IGMP specifications 15
6.3.3.1 Additional state "reports reception disable" 15
6.3.3.2 Modified timer value 15
6.3.3.2.1 Query response interval 15
6.3.3.3 Modified actions 15
6.3.3.3.1 "Start timer*" action 15
6.3.3.3.2 "Start retransmit timer" action 15
6.3.3.4 Additional actions 16
6.3.3.4.1 "Resend report" 16
6.3.3.4.2 "Send group specific query (max eff)" 16
6.3.3.4.3 "Send group specific query(0,1)" 17
6.3.3.4.4 "N member +" 17
6.3.3.4.5 "N Member -" 17
6.3.3.4.6 "N Member 0" 17
6.3.3.4.7 "Start disable_timer" 17

Annex A (informative): Description of S-IGMP 18

A.1 Outline of IGMP 18
A.1.1 Dynamic vs. static routing 18

A.2 IGMP issues in a satellite system 18
A.2.1 IGMP message flooding 19
A.2.1.1 Numerical examples of adaptation impact 19
A.2.1.2 Impact of limiting the retransmitted report burst 20

A.2.2 Latency of stopping multicast transmission 21
A.2.2.1 Adaptation solution for latency 21
A.2.2.1.1 Max_Response_Time tuning....................21
A.2.2.1.2 Rexmt tuning....................21

Annex B (informative): Description of scenarios for alternative versions of IGMP 22

B.1 IGMP querier interface modes and options....................22

Annex C (informative): Bibliography....................23

History 24

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (http://webapp.etsi.org/IPR/home.asp).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Technical Specification (TS) has been produced by ETSI Technical Committee Satellite Earth Stations and Systems (SES).

Introduction

The general Service and Architecture aspects for Broadband Satellite Multimedia (BSM) systems are described in TR 101 984, IP-over-satellite aspects are described in TR 101 985 and functional models are defined for Quality of Service (QoS), Addressing, Routing and Multicasting. TR 102 156 develop the requirements for multicasting in further detail.

The present document continues this theme by addressing an efficient solution for IPv4 multicast group management over BSM systems based on the definitions in the TR 101 984, TR 101 985 and TR 102 156.

The specification of IGMP adaptation herein is based on results of the GEOCAST project of the EC's IST Programme. The IP multicast-over-satellite scenario described herein is also under consideration for the Amerhis project.

1 Scope

The Internet Group Management Protocol (IGMP) is an integral part of IP and is required to be implemented by all hosts wishing to receive IP multicasts.

The present document specifies modifications to IGMP which improve its performance when applied over satellite links. These modifications are internal to IGMP and do not affect the system's interoperability with IP. The modifications concern particularly the so-called "IGMPv2 mode" (described in RFC 2236 [2]) which falls under the umbrella of IGMPv3 (RFC 3376 [1]).

IGMPv3 offers automatic backwards reversion capability with the alternative IGMP modes.

The adapted mode of IGMP is termed S-IGMP.

This solution is suitable for multicast groups on geostationary satellite systems with two-way links. It is suitable for systems based on Any-Source Multicast (e.g. star networks) and especially for those which do not automatically retransmit IP multicast packets on return links to all group members. It is also beneficial to systems which do have this ability for full subnet multicast.

2 References

The following documents contain provisions which, through reference in this text, constitute provisions of the present document.

- References are either specific (identified by date of publication and/or edition number or version number) or non-specific.
- For a specific reference, subsequent revisions do not apply.
- For a non-specific reference, the latest version applies.

Referenced documents which are not found to be publicly available in the expected location might be found at http://docbox.etsi.org/Reference.

[1] IETF RFC 3376: "Internet Group Management Protocol version 3" (IGMPv3).

[2] IETF RFC 2236: "Internet Group Management Protocol version 2" (IGMPv2).

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

multicast group: multicast IP address to which hosts may subscribe

proxy: function that intervenes between a source and destination and performs that function as an intermediary for the remote devices in each direction

NOTE: For multicast group management an IGMP Proxy behaves as a single IGMP client on behalf of several downstream hosts, and in the opposite direction as a local IGMP querier to these hosts on behalf of a remote querier.

snooping: function associated with a layer 2 switch or bridge that intervenes on a given layer 3 protocol between a source and destination by learning about network behaviour from intercepted IP packets and without explicit configuration as a network function (with IP address, etc.)

3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

ASM	Any-Source Multicast
BSM	Broadband Satellite Multimedia
GES	Ground Earth Station
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IPv4/v6	Internet Protocol version 4/6
LAN	Local Area Network
PIM	Protocol Independent Multicast
QoS	Quality of Service
RFC	Request For Comments
SDAF	Satellite Dependent Adaptation Function
SI	Satellite Independent
SIAF	Satellite Independent Adaptation Function
SSM	Source-Specific Multicast
ST	Satellite Terminal
STF	Special Task Force
UES	User Earth Station

4 Overview and summary

4.1 Background

Clause 4 provides a general description of the background and choices for IGMP adaptation.

The BSM system is intended to interwork seamlessly with IP networks. In the case of IP multicast, group management protocols for local hosts on the satellite segment need to be interfaced with the relevant routing protocols on the external part of the network, to ensure dynamic configuration of multicast routes.

The present document defines how IGMP needs to be adapted to suit the BSM environment, whilst maintaining the same external interface and behaviour towards IP. The adapted mode of IGMP will be termed S-IGMP.

The BSM network and protocol architecture relies on the generic BSM architectures in TR 101 985. The background to BSM IP multicast interworking is described in TR 102 156.

Applicable BSM IGMP multicast architectures are described in TS 102 294.

4.2 Need for IGMP adaptation

IGMP is used as the basis for IP multicast forwarding across a local network attached to a multicast router. IGMP is intended particularly for LAN's with a shared (broadcast) medium, where every host can listen to IGMP reports sent by others, and where there is low delay.

As satellite multicast groups become very large and dynamic, serious scalability consequences can arise. This is particularly true of typical satellite configurations where, even with two-way communications, the broadcast property generally exists only in the forward link; terminals cannot listen to direct or retransmitted reports from other terminals to suppress redundant reports. The IGMP querier located at the gateway router may then receive thousands of reports from the ST's (hence the name "IGMP flooding") leading to a waste of bandwidth as well as high processing power demand at the querier.

Another problem is significant waste of traffic resources due to latency in stopping multicast group transmissions when members have left. The delay introduced in hosts' reports to improve report suppression further degrades this latency. A mechanism to adjust the host delay for optimal latency without significantly affecting report suppression is defined.

A more detailed explanation of the need for and choices of adaptation is given in clause A.1.

4.3 IGMP modes

IGMPv3 (RFC 3376 [1]) is the currently avalaible version of the protocol, but it can be configured, or can reconfigure itself automatically depending on its IGMP client versions for each group and interface, to operate in one of:

1) IGMPv3 native mode:
   - intended for interoperability with any-source multicast (ASM), with extensions for source-specific multicast routing protocols (e.g. PIM-SSM);
   - has scalability problems over large satellite multicast groups.
2) IGMPv2 mode:
   - intended for interoperability with a wide range of ASM routing protocols (e.g. PIM-SM).
3) IGMPv1 (RFC 1112) mode, for legacy applications.

S-IGMP described herein concerns modifications to IGMPv2, or equivalently to the "IGMPv2 mode" of IGMPv3. This version is chosen because it:

- is the most efficient IGMP mode in terms of signalling traffic (allows report suppression);
- needs a minimum of adaptation for satellite use compared to IGMPv3 mode.

4.4 Advantages of S-IGMP solution

The S-IGMP solution defined herein is designed for high performance and scalability consistent with:

1) High reuse of existing IGMP functions:
   - the IGMP client functions are unchanged (existing IGMPv2 end-hosts and proxies, or this mode in IGMPv3 clients can be used);
   - the IGMP packets on the satellite link are fully consistent with IGMPv3 (RFC 3376 [1]).
2) Low implementation complexity:
   - only the IGMP querier function (in the router) is modified.

A key feature of the adaptation is that the reuse of the existing IGMP packet format and standard IGMPv2 clients simplifies deployment of STs and user equipment. Only the behaviour of the querier is modified.

Any adaptation of IGMPv3 mode with similar performance would require modified clients as well as queriers.

4.5 Influence of IGMP proxies, etc.

Like many other IP functions, IGMP is often handled through proxy servers (or similar intermediaries) when a single interface to a local subnetwork is needed. Such IGMP proxies are common in the marketplace and require no modification if they are used at the client side of S-IGMP. (If they are used at the querier side of S-IGMP then they must implement S-IGMP instead of the source router). The present document does not therefore define the functions of such proxies (or snooping switches), but it describes their use in network architectures.

5 Applicability and assumptions

5.1 Use of IGMP in BSM multicast architectures

A common scenario for IP multicast over satellites is the so-called "edge" multicast network, in which only single hosts or small networks of hosts (i.e. "stub" networks with no transit traffic) are attached to the remote STs and are provided with multicast services via satellite from the Internet or from sources situated in other networks.

In this scenario a multicast router situated on one side of a satellite system needs to forward IP multicast packets to the remote IP hosts (or multicast "receivers"). At the IP network level, this situation is very similar to the terrestrial multicast case concerning the multicast router's forwarding on an interface to its locally attached hosts.

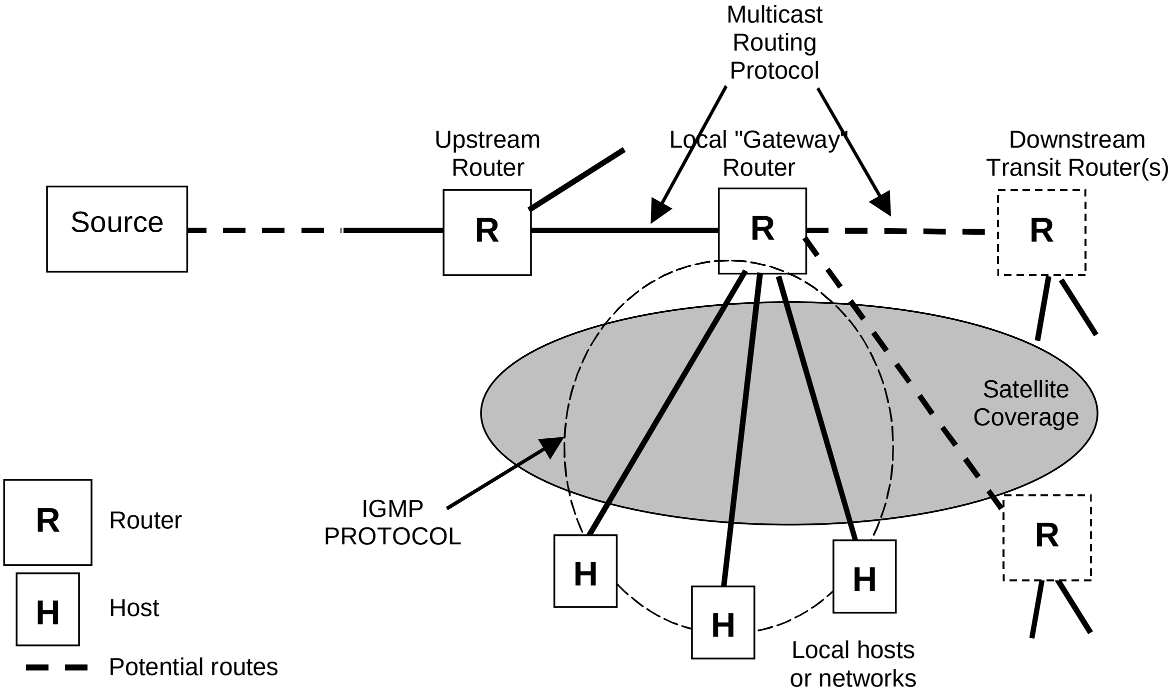

Figure 1: Typical BSM edge IP multicast network scenario

The "gateway" multicast router may in general have one or more interfaces, as indicated, to local hosts as well as to other routers, via both terrestrial and satellite links.

The BSM system internal protocol that is needed to maintain the dynamic multicast forwarding table in the "gateway" router towards local hosts could be chosen arbitrarily by the satellite system operator, provided that the BSM system interworks with external protocols. However it is logical and simpler to base this protocol on IGMP (as in IPv4 networks elsewhere), not only for directly attached hosts but also in the more typical case where the hosts are on local networks and are connected to the gateway router via a layer 3 (IGMP proxy) or layer 2 (snooper) LAN device.

This IGMP-based solution has the following advantages over full routing protocol implementation (e.g. via LAN multicast routers) over the BSM:

1) simpler to implement, above all in STs (for lower cost). Only active groups on the air interface of the gateway router need to be maintained instead of complete routing tables;

2) independent of, and compatible with most routing protocols from the source;

3) more scalable.

5.2 IGMP modes and multicast routing protocol options

IGMPv2, the mode chosen as the basis for S-IGMP, has been officially made obsolete by IGMPv3. To permit interworking with earlier versions, however, IGMPv3 allows reversion to IGMPv2 mode. The adaptation of IGMPv2 for satellites discussed herein therefore can be considered as this mode of IGMPv3, which is identical to the original IGMPv2.

IGMPv3 mode cannot be easily adapted for efficient use over satellites as it forbids suppression of redundant responses from hosts, and more complex modifications not only to the querier but also to the clients would be required to achieve the same.

A router must have an instance of IGMP operating on each relevant interface, and each instance could be a different version of IGMP.

IGMPv2 is intended to interwork with ASM routing protocols (including PIM-SM for example). IGMPv3 adds support for multicast "source filtering" in order to interwork with PIM-SSM (source-specific multicast) which allows routing from multiple sources rather than via a single rendezvous point as in PIM-SM, for example. PIM-SSM however is currently less often implemented.

The architectural scenarios for and selection between alternative versions of IGMP are described in clause A.2.

5.3 IGMP network architecture applicable to BSM

In the edge multicast scenario two main network architectures for the use of IGMP are conceivable, according to whether the end-host connects to the BSM system via:

1) Layer 3 STs (acting as IGMP proxies); or

2) Layer 2 STs (switches/bridges with optional "transparent" layer 3 functions e.g. snoopers).

These, and more general, cases are described in TS 102 294.

In the above cases, standard commercial IGMP proxies or snooping switches can be used.

5.4 IGMP-specific BSM functional model

The general IGMP protocol architecture applicable to BSM is given in TS 102 294.

Figure 2 shows a more detailed relationship between IGMP over the satellite and the control/management plane protocol stacks in the BSM system.

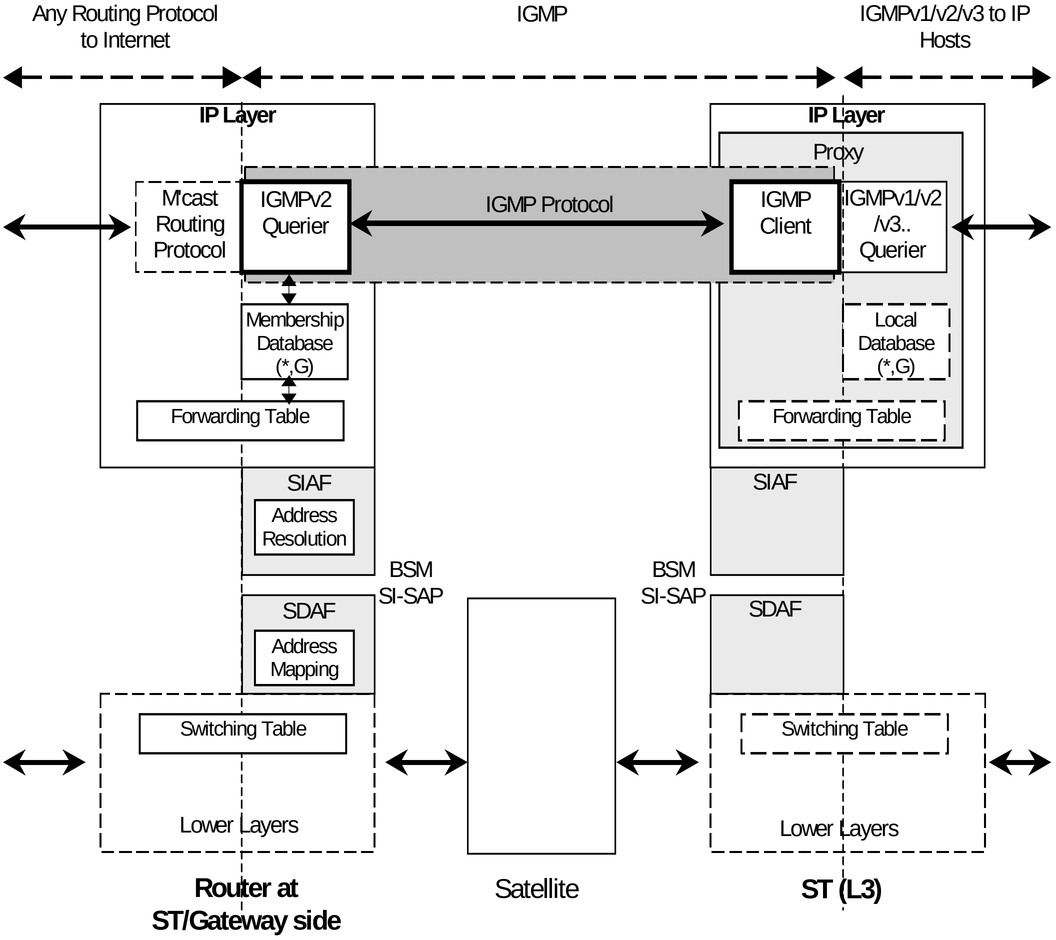

Figure 2: IGMP protocol architecture over the BSM

Figure 2 illustrates in particular that the main result of the IGMP protocol is to establish the membership database in the querier part of the router of the form (*,G) indicating that there is at least one member of the interface interested in receiving a multicast address G and irrespective of the source (*).

An example of an IGMPv2 group membership table for one IGMP router interface, combined with forwarding and address resolution data, is as follows.

Table 1: Example of combined multicast group membership/forwarding table

Multicast group	IP Host @	Host interface	IGMP version	Reception time last host report	Group timer value	Group rexmt timer
234.1.1.1	184.10.1.1	ST@n	2	10:34:15	Tx(x)	Txr(a)
	184.10.1.2	ST@m	2	10:34:12		
234.1.1.2	184.10.1.1	ST@n	2	09:31:02	Ty(y)	Tyr(b)
	184.10.1.3	ST@p	2	08:01:07		

5.5 IGMP issues in a satellite system

IGMP in general raises the following performance issues:

flooding: This occurs when many IGMP clients reply to a broadcast request from the IGMP querier at the same time, flooding it with report messages. It may also occur in the other direction if these messages are retransmitted by the gateway station to all members of the group, by an automatic re-multicast mechanism.

latency: This is the delay for stopping a multicast transmission after the last client leaves a multicast group. It is the delay of the querier becoming aware that the group is empty. Latency is a consequence of the anti-flooding mechanism.

The IGMP protocol has been designed for terrestrial networks with shared many-to-many media (e.g. Ethernet) and with low transmission delay between nodes. When IGMP is applied to satellite systems, increased flooding and latency and hence waste of valuable resources can result. Adaptation of IGMP to the technical characteristics of satellite links is therefore desirable.

In order to resolve flooding and latency issues, the IGMP adaptation is based on two techniques:

- limited retransmission of received reports over the air interface enabling emulation of a shared LAN medium, sufficient to ensure report suppression;
- tuning of IGMP Protocol parameters, particularly the Max_Response_Time of reports.

These modifications apply to IGMPv2 mode and only to the querier, and thus imply minimum system implementation or modification.

5.6 IGMP messages

No change to IGMPv2 message structure is specified, and only IGMP behaviour is modified. To recapitulate, IGMP messages are encapsulated in IP datagrams, with an IP protocol number of 2. All IGMP messages described in RFC 2236 [2] and in the present document are sent with an IP TTL of "1", and should contain the IP Router Alert option (see RFC 2113) in their IP header.

6 S-IGMP specification

6.1 General

The present document defines modifications to the Internet Group Management Protocol IGMPv2 (see RFC 2236 [2]), or more specifically to the IGMPv3 protocol (see RFC 3376 [1]) in its IGMPv2 mode of operation. The modified version of IGMP is referred to here as S-IGMP.

S-IGMP applies primarily to the modification of the behaviour of the IGMP querier located in the local router, so that only the querier function needs to be changed. The modifications to the querier function consist of:

1) introduction of a new state in the state diagram;

2) setting of specific timer values;

3) modified actions;

4) introduction of new actions in the state diagram.

NOTE: In setting the timer values below, care must be taken to ensure they are consistent among other IGMP routers on the same link.

6.2 S-IGMP client

The IGMP client shall function as specified in RFC 2236 [2] section 6 (as referred to in RFC 3376 [1]). It shall process one state diagram per multicast group active in the membership database.

To clarify the implementation required, and since RFC 2236 [2] leaves some options open, the further requirements below clarify or specify items required for satisfactory adaptation performance:

1) The unsolicited report interval of the IGMP Client shall be configurable.

2) In addition to the specified actions for the "join group" event, the IGMP client shall not send an unsolicited report if it receives a (retransmitted) report before the unsolicited report timer expires.

6.3 S-IGMP router (querier)

The S-IGMP querier function is an adaptation of the IGMPv2 querier of RFC 2236 [2], and is located in the local router. This clause addresses the state diagram of the querier.

6.3.1 Overall querier state diagram

The "querier" mode of the router shall function as per the state diagram specified in RFC 2236 [2] section 7, as shown in figure 3.

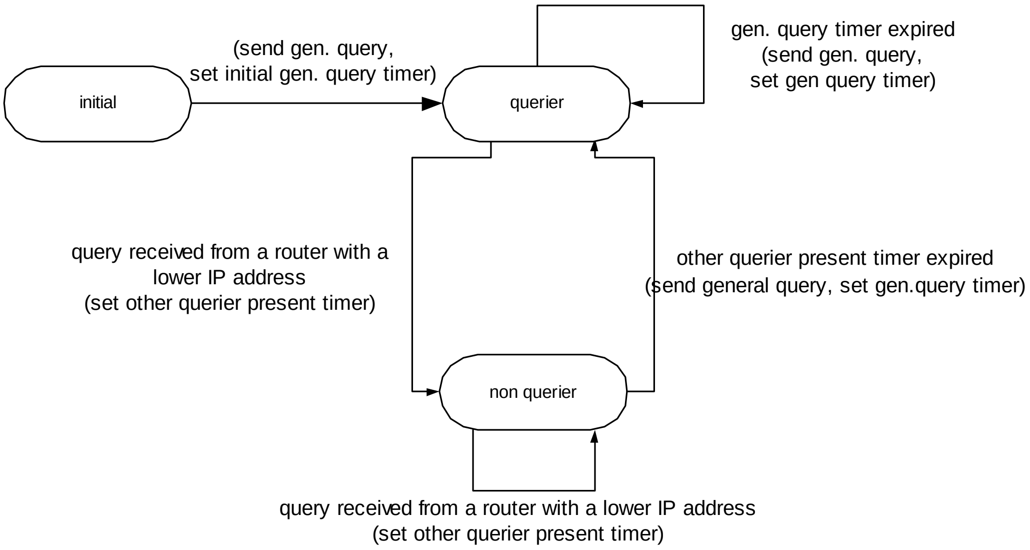
initial
(send gen. query,
set initial gen. query timer)
querier
gen. query timer expired
(send gen. query,
set gen query timer)
query received from a router with a
lower IP address
(set other querier present timer)
other querier present timer expired
(send general query, set gen.query timer)
non querier
query received from a router with a lower IP address
(set other querier present timer)

Figure 3: IGMP querier general state diagram

Figure 3 indicates that the downstream interface of the router can act either as an IGMP querier or a non-querier (e.g. conventional router).

For S-IGMP the following timer values shall be set within the above state diagram:

1) Startup query interval.

[startup query interval] shall be set to [query interval] (instead of the default of 0,25 × [query interval]). This occurs during the **"set initial gen. query timer"** action, and is chosen to avoid overlapping sets of query responses due to the response delays which could occur for the default value.

2) Query response interval.

[Query Response Interval] shall be set to 25,5 s (i.e. the maximum permitted, instead of the IGMP default value of 10 s). This corresponds to the Max_Response_Time of the general query message sent to the all-systems group (224.0.0.1), and is set in order to maximally spread the responses.

6.3.2 Router state diagram in querier mode

As shown in figure 4, the S-IGMP querier shall be in one of four states for any IP multicast group on the interface.

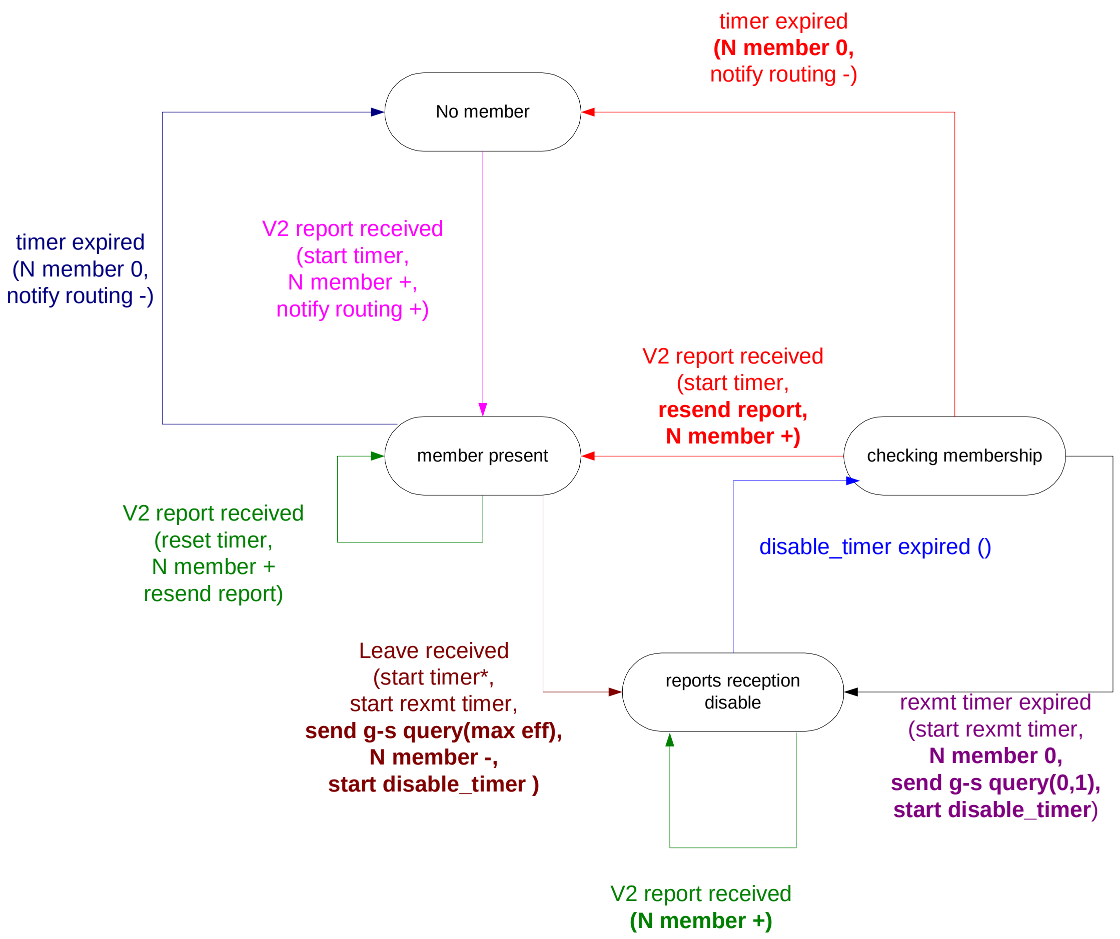

Figure 4: S-IGMP state diagram per member group in IGMPv2 querier mode

One such diagram applies to each multicast group on the interface. The state and actions corresponding to reversion to IGMPv1 mode are not shown, since they are identical to those in RFC 2236 [2].

Compared to RFC 2236 [2] section 7, the following changes have been made:

1) an additional state "Reports reception disable" has been created;

2) specific timer values have been set;

3) modified actions;

4) the following new actions have been created (indicated in bold on figure 4):

- resend report;
- send group specific query(max eff);
- send group specific query(0.1);
- N member +;
- N member -;
- N member 0;
- start disable_timer.

6.3.2.1 Parameter definitions

The following parameter definitions have been introduced for use in the timer and action values in table 2.

Table 2

Parameter	Unit	Definition
RTT	s	Round Trip Time – the time delay of a signal sent from one ST to another and back to the ST over the satellite. The IGMP adaptation defined here is nevertheless based on the use of geostationary satellites.
Forward Robustness		the minimum number of reports retransmitted per query, chosen to counter losses
P		parameter (between 0 and 1, default: 0,95) to reduce the latency of stopping data forwarding. It corresponds to the probability of receiving a report to the first group-specific query.
Remove_time_IGMPv1	s	time before removing an IGMPv1 host from a membership list.
Remove_time_IGMPv2	s	time before removing an IGMPv2 host from a membership list.
N		the estimated number of members per multicast group, and per interface, calculated by the querier.
k		the count of retransmitted reports following a query.
NOTE: The number N does not need to be estimated with any significant accuracy, and a sufficient approximation can be deduced from the number of "join" messages received.		

6.3.3 Detailed S-IGMP specifications

6.3.3.1 Additional state "reports reception disable"

In this state the IGMP querier shall ignore membership reports for this specific group received during a RTT period after sending a group specific query. Reports received during this period could arrive in response to a previous group-specific query or a general query.

6.3.3.2 Modified timer value

6.3.3.2.1 Query response interval

As in the overall querier state diagram in clause 6.3.1, [Query Response Interval] shall be set to 25,5 s (i.e. the maximum permitted, instead of the default value of 10 s). This corresponds to the Max_Response_Time of the general query message sent to the all-systems group (224.0.0.1), and is set in order to maximally spread the responses and thus to maximize suppression of unnecessary responses.

6.3.3.3 Modified actions

6.3.3.3.1 "Start timer*" action

The action of "start timer*", for the multicast group timer on the interface, shall be as follows:

set the timer to [(RTT+MAX(Max_Response_Time.(1- $\sqrt[N]{1-P}$), 0,1))+(Robustness-1).(RTT+0,1)].

6.3.3.3.2 "Start retransmit timer" action

The action of "Start Retransmit Timer", for the multicast group "rexmt timer" on the interface, shall be as follows:

if N ≠ 0:	set the timer to:	[RTT+MAX(Maximum_Response_Time.(1- $\sqrt[N]{1-P}$), 0,1)],
else:	set the timer to:	RTT+0,1.

In the above actions, Max_Response_Time is optimized to reduce the report flooding and the latency (values are assigned depending on the number of members assumed).

6.3.3.4 Additional actions

The use of the actions described below is as shown in the State Diagram.

6.3.3.4.1 "Resend report"

This consists of retransmitting the membership reports under certain conditions.

A Membership Report shall be sent over the air interface only if the parameter k is lower than Forward_Robustness.

Moreover, if the delay between a membership report and the previous one is lower or higher than max_delay, the counter shall be increased or reset respectively.

The max delay shall be chosen as a function of Forward_robustness as follows:

$$\text{Max delay} = \frac{3\times 25{,}5}{(25{,}5\times(forward\_robustness-1)/RTT)-1}$$

Examples of max delay (for RTT = 0,6 s):

Forward robustness	Max delay
3	0,889 s
4	0,593 s
5	0,447 s
6	0,357 s

6.3.3.4.2 "Send group specific query (max eff)"

The group specific query shall be sent on the air interface with a Max_Response_Time as follows:

If N =0;

Then set Max_Response_Time = 0,1 s.

$$\text{Else if N} < \text{Round.sup}\left(\frac{-1+\sqrt{1+4\times\left(\frac{25{,}5}{RTT}\right)^2}}{2}\right);$$

Then set maximum response time = round($\sqrt{N\times(N+1)}\times RTT$) s

$$\text{Else if N} \geq \text{Round.sup}\left(\frac{-1+\sqrt{1+4\times\left(\frac{25{,}5}{RTT}\right)^2}}{2}\right);$$

Then set maximum response time = 25,5 s, set k = 0.

In the above equations:

- "Round.Sup" shall round up the value to the next integer.
- "Round" shall round down the value with a precision of 1/10.

6.3.3.4.3 "Send group specific query(0,1)"

The group specific query is sent on the air interface with a Max_Response_Time set to 0,1 s. Moreover, the counter parameter k shall be reset.

6.3.3.4.4 "N member +"

This action shall update the membership list following an IGMP Membership Report as follows:

- If the member is already in the membership table, its last membership table entry shall be updated.
- Else:

1) If the number of table entries is less than $\text{Round.sup}\left(1{,}5\times\left(\frac{-1+\sqrt{1+4\times\left(\frac{25{,}5}{RTT}\right)^2}}{2}\right)\right)$,

   and if the member is new, the member shall be added to the membership.

2) If the number of table entries is more than $\text{Round.sup}\left(1{,}5\times\left(\frac{-1+\sqrt{1+4\times\left(\frac{25{,}5}{RTT}\right)^2}}{2}\right)\right)$,

   the member shall be added by removing the oldest from the membership list.

Moreover, for all the other lines in the membership list, the time of the last received Membership Report shall be compared to a value V given below. If the time is greater than V, the entry in the membership list shall be deleted.

- For IGMPv1 hosts, V= [current time - Remove_time_IGMPv1].
- For IGMPv2 hosts, V= [current time - Remove_time_IGMPv2].

6.3.3.4.5 "N Member -"

This action updates the membership list by removing the entry for the member reporting a "leave".

6.3.3.4.6 "N Member 0"

This action resets the membership list to zero.

6.3.3.4.7 "Start disable_timer"

This action sets the timer to RTT.

Annex A (informative): Description of S-IGMP

The following text is offered for information only.

A.1 Outline of IGMP

IGMP is designed to support the forwarding (and, more generally, routing) of IP multicast packets from a router to its locally attached hosts.

IGMP is an integral part of IPv4 and is required to be implemented by all hosts wishing to receive IP multicasts.

IGMP enables the router to maintain a database of the multicast groups of its connected hosts. The router can then forward multicast packets to hosts as needed, and also use the database to participate externally in multicast routing with other routers.

IGMP consists of two main parts:

1) The "querier" is used by the local router to maintain a multicast group membership database of its connected hosts.

2) The "client" is used by IP hosts to report their multicast group memberships to the querier.

A.1.1 Dynamic vs. static routing

In multicasting, the forwarding of IP multicast packets by a router to the multicast groups is based on a forwarding table, which links multicast IP addresses to local links and to next hop IP addresses. The complexity of this table depends on the satellite system and network architecture (e.g. allocation of groups of receivers to satellite broadcast physical channels i.e. antenna beams, TDM time slots, frequency channels, etc.).

A simple solution is to configure this table statically (manually), so that most receivers are flooded with traffic to guarantee reception. This may be wasteful of satellite resources, and a better solution is often dynamic multicast where the forwarding table is maintained according to the demand from users.

IGMP is intended to support dynamic multicast routing in the BSM system.

A.2 IGMP issues in a satellite system

IGMP in general raises the following performance issues:

flooding: this occurs when many IGMP clients reply to a broadcast request from the IGMP querier at the same time, flooding it with report messages which are not suppressed to the normal extent due to the satellite system characteristics.

latency: this is the delay for stopping a multicast transmission after the last client leaves a multicast group. It is the delay of the querier becoming aware that the group is empty in order to stop multicast forwarding. Latency is a consequence of the anti-flooding mechanism.

The IGMP protocol has been designed for terrestrial networks with shared many-to-many media (e.g. Ethernet) and with low transmission delay between nodes. When applied to satellite systems IGMP can lead to inefficient use of valuable resources, and adaptation of IGMP to the technical characteristics of satellite links is therefore desirable.

The solutions to be specified are based on tuning of IGMP parameters.

A.2.1 IGMP message flooding

The IGMP protocol is well-suited to a terrestrial network and shared media, in which every host can listen to reports sent by others. If a report has been sent from another host, the host will then stop its timer and not send its report, to prevent sending useless reports.

However, this is not the case in a satellite context; terminals typically cannot listen directly to replies from other terminals. Moreover, as satellite multicast groups can be very large (problem of scalability) and very dynamic, this can have serious consequences. The IGMP querier located at the Hub may receive thousands of reports from the ST (hence the name "IGMP flooding") leading to a waste of bandwidth and high CPU demand at the querier side. (N.B this is different from the term often associated with flooding, or retransmitting, by an IP host of these reports to all interfaces)

As an example, consider a single multicast group and a satellite system with 5 000 satellite terminals with at least one member of this group behind each of them. If IGMP is implemented as defined in RFC 2236 [2] (without adaptation and without retransmission of reports at the Hub), at each query the system will have to transmit (5 000 × 32 × 8) **1,28 M of return bits** of IP packets from the satellite terminals (UES) to the Gateway (GES).

Figure A.1 indicates why IGMP needs to be adapted to the satellite context in a transparent way for hosts.

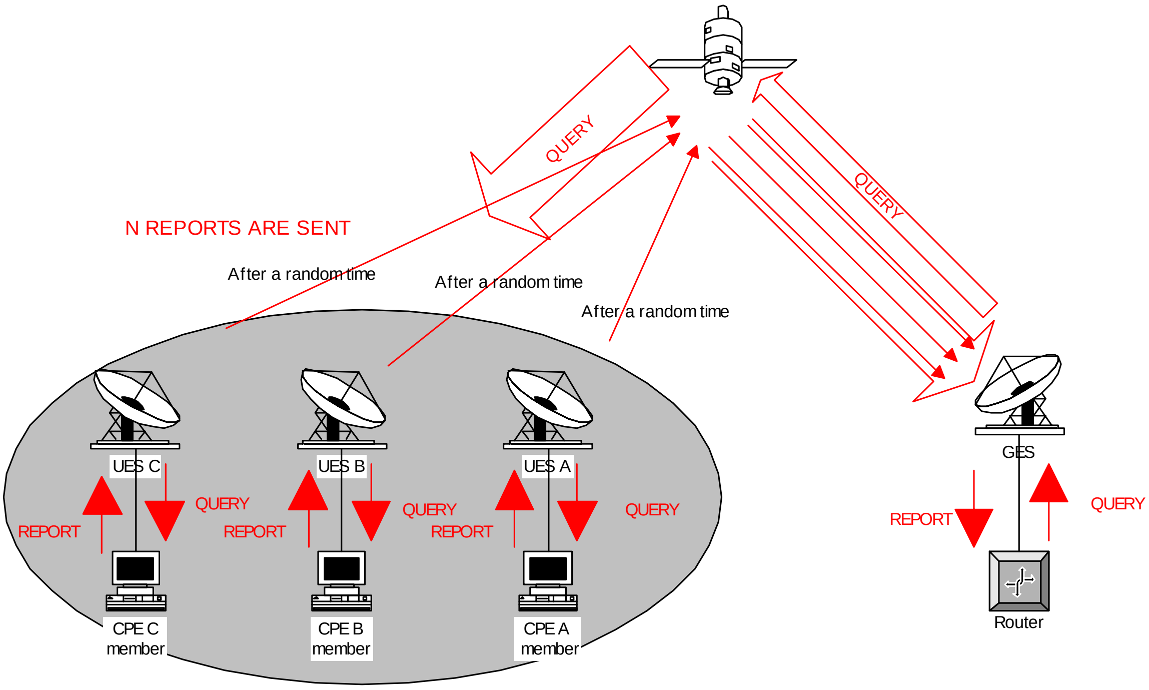

Figure A.1: IGMP flooding in the satellite context

It is this "flooding" of the reports on the return link that creates scalability problems, since they each consume satellite resources and processing power in the Hub equipment, in contrast to the small number of query messages on the forward link.

A.2.1.1 Numerical examples of adaptation impact

The total number of IGMP reports transmitted over the satellite with or without retransmission of reports by the querier and with different values of Max_Response_Time is as follows.

Table A.1: Mean number of reports / query / group

Number of members		**1**	**10**	**50**	**100**	**500**	**1 000**	**5 000**
		total number of reports						
Without adaptation (no rebroadcast)	Max_Response_Time = 0,1 s	1	10	50	100	500	1 000	5 000
	Max_Response_Time = 25,5 s	1	10	50	100	500	1 000	5 000
With Adaptation	Max_Response_Time = 0,1 s	1	10	50	100	500	1 000	5 000
	Max_Response_Time = 25,5 s	1	1,23	2,17	3,35	12,7	24,5	118

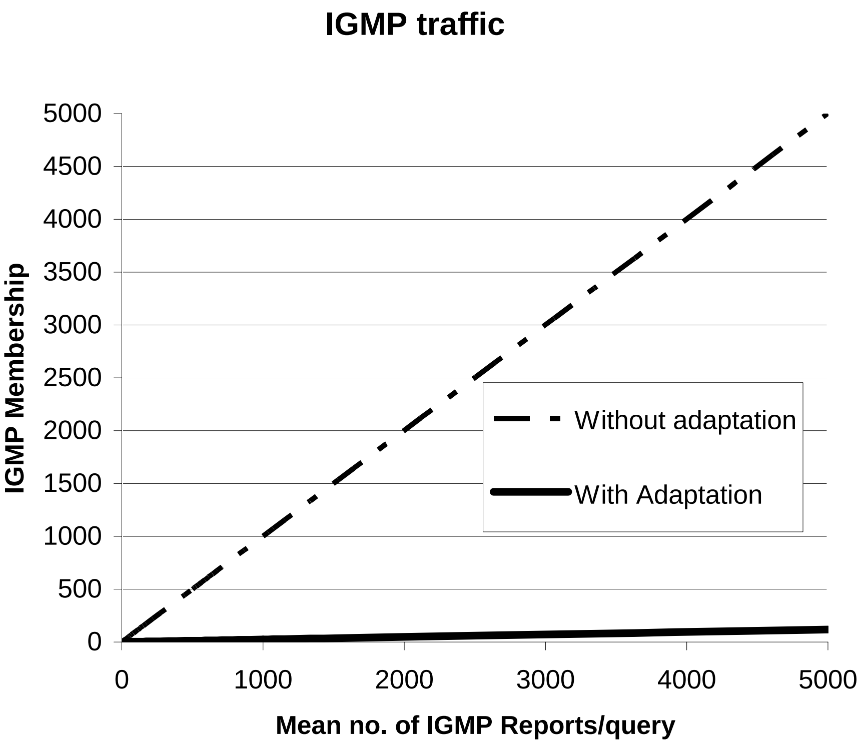

Figure A.2: Adaptation performance - retransmission on the air interface (Max_Response_Time = 25,5 s)

Note that the figure A.2 and table A.1 deal with the both the general queries and group specific queries.

The maximum IGMPv2 value of 25,5 s for Max_Response_Time is used to maintain compatibility with standard IGMPv3 clients, and to spread maximally the responses over time in order to allow maximum suppression.

While the Max_Response_Time for a general query should always be set to the maximum value 25,5 s, for a group specific query the Max_Response_Time should be optimized as specified in the previous parts.

It can be seen that the benefit of the adaptation is greater for the large groups: instead of transmitting 5 000 reports, only 118 will be transmitted leading to a reduction of 97 %, and a large bandwidth saving.

For a group of 5 000 members, with a Max_Response_Time of 0,1 s, the terminals will need 5 000 × 32 × 8 = 1,2 Mbits to transmit the reports. For the same group, with a Max_Response_Time of 25,5 s, the terminals will need 118 × 32 × 8 = 30 kbits to transmit reports.

A.2.1.2 Impact of limiting the retransmitted report burst

The number of reports retransmitted by the IGMP querier is limited to [Forward Robustness] using an algorithm to detect the start of the burst.

As shown in clause A.1.2.1, with a multicast group of 5 000 members and a RTT of 0,6 s, the adaptation results in typically 118 reports on the return channel. Of these the adaptation further limits rebroadcast by the gateway of only the first 3 or 4 (sufficient to ensure report suppression) reports (as defined by the "resend report" action), which would in this case save 29 kbits on the forward link. This represents a further gain of efficiency.

A.2.2 Latency of stopping multicast transmission

Another important issue of IGMP behaviour is the latency in stopping transmission after the last host leaves a multicast group.

In order to decrease effects of flooding as we have seen above, the period over which reports are sent is increased by setting the Max_Response_Time parameter.

But when the last member is leaving, IGMP implemented as per the RFC will take time, depending on the report spreading period, to detect there are no members in the group and stop transmitting the multicast stream on the air interface. The time required is defined by the IGMP protocol and will lead to a waste of bandwidth. There is a trade-off between spreading the flooding and time needed to stop the flow.

As an example, consider that the "Max_Response_Time" (parameter of the group queries) is set to 25,5 s in order to spread the flooding. Also consider that the last member of Multicast 1Mbps video streaming group is leaving. It detects there is no member left in this group, and the querier needs about 51 s (Robustness = 2 (IGMPv2 default value)); this costs about **51 Mbits on the forward link**.

A.2.2.1 Adaptation solution for latency

As shown above there is a conflict between IGMP optimization for flooding and for latency, and a trade-off or more sophisticated optimization needs to be found.

The solution is tune two IGMP parameters during the leave process:

1) Max_Response_Time is varied by group-specific query messages.

2) The "Rexmt" timer sets the delay between the first and the subsequent group-specific queries.

This solution has minimal impact on the report suppression to minimize flooding. The effect on the graph of figure A.2 is not significant as only a few additional reports in the region from 0 to 10 members are unsuppressed.

A.2.2.1.1 Max_Response_Time tuning

For the first group-specific query after reception of an explicit leave message, the Max_Response_Time is set according to the estimated number of members in the group, N. This timer value is optimized to ensure that few reports are received if there are still members and for latency to leave if there was only one member remaining. Hence there is minimal impact on flooding improvement.

For subsequent group-specific queries (if there is no report at the first g-s query), the max response time is set to the minimum, for minimum latency.

A.2.2.1.2 Rexmt tuning

The "**Rexmt**" timer set up by the querier between 0 and [max_response_time set by first g-s query + RTT] seconds.

The Rexmt is adapted to the satellite context and should be computed according to the estimated group size and the parameter P.

Annex B (informative): Description of scenarios for alternative versions of IGMP

B.1 IGMP querier interface modes and options

There are several options available for the type of IGMP available over the BSM network depending on the type of host clients and proxies used. The following IGMP network scenario is used as a basis for discussion below.

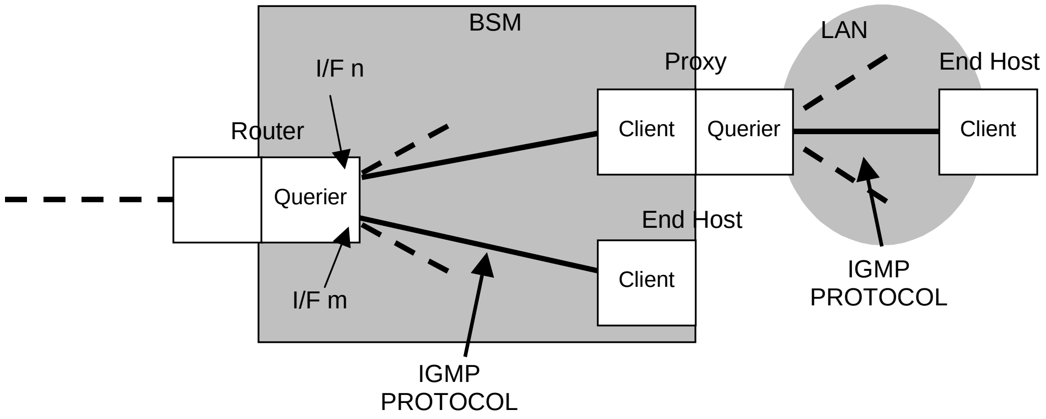

Figure B.1: IGMP scenario

For each IGMP querier interface the lowest version of all attached IGMP client may be used to determine the version of IGMP in operation, as follows.

Table B.1: Permissible IGMP modes showing S-IGMP applicability

	Querier Interface mode - Router	Client type or mode - Proxy	Querier mode - Proxy	Client type or mode - Host	Comments
	With proxy				
	V1	V1	V1	V1	e.g. Legacy LANs
S-IGMP Modes	V2	V2	V1	V1	e.g. Legacy LANs
	V2	V2	V2	V2	
	V2	V2	V3	V3	
	V3	V3	V3	V3	For full PIM-SSM etc. in network
	Without proxy				
	V1			V1	e.g. Legacy LANs
S-IGMP Mode	V2			V2	
	V3			V3	For full PIM-SSM etc. in network

The table indicates that in this network scenario, S-IGMP will be automatically enabled only if there is at least one IGMPv2 client either in a proxy or in a directly connected end host, at each router querier interface.

Alternatively the IGMPv3 queriers can be explicitly configured to operate in S-IGMP (or IGMPv2) mode, and any IGMPv3 clients would automatically switch to IGMPv2 mode.

Annex C (informative): Bibliography

- ETSI TR 101 985: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; IP over Satellite".
- ETSI TR 102 156: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; IP interworking over satellite; Multicasting".
- ETSI TS 102 294: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM) services and architectures; IP interworking via satellite; Multicast functional architecture".
- ETSI TR 101 984: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; Services and Architectures".
- IETF RFC 1112: "Host Extensions for IP Multicasting" (IGMPv1).
- IETF RFC 2113: "IP Router Alert Option".

History

Document history		
V1.1.1	February 2004	Publication