ETSI TS 102 674 V1.1.1 (2009-11)

Technical Specification

**Satellite Earth Stations and Systems (SES);
Broadband Satellite Multimedia (BSM);
PIM-SM Adaptation**

Reference

DTS/SES-00291

Keywords

broadband, interworking, IP, 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, please send your comment to one of the following services:
http://portal.etsi.org/chaircor/ETSI_support.asp

Contents

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

PIM-SM is the mode of the PIM protocol most widely used in existing and proposed multicast routing applications today. In the following discussion "PIM-SM" is understood to mean the messages and operation of this protocol.

The way in which PIM-SM may be carried over satellite is described herein. In particular, ways in which PIM-SM can be adapted over satellite links are defined, which may lead to more efficient use of satellite resources.

1 Scope

The present document specifies the way in which PIM-SM multicast routing protocols may be used over satellite systems. From the many possible scenarios for PIM-SM implementation over different network configurations, three types of PIM scenarios are initially considered:

- Native PIM-over-Satellite;
- Adapted PIM-over-Satellite (S-PIM);
- PIM Snooping/Proxying.

In particular, an adaptation of the PIM-SM standard to a non-standard BSM-internal version of PIM (i.e. S-PIM) is described. In addition, PIM-SM protocol configuration, and/or proxying needed to enable efficient operation over the satellite system and the associated interworking with terrestrial networks is described.

The present document builds upon previous BSM documents referenced in annex B, and notably:

- TS 102 293 [4]: Multicast Group Management; IGMP adaptation.
- TS 102 294 [1]: BSM Multicast Functional Architecture.
- TS 102 461 [i.2]: Multicast Source Management.

The PIM-SM protocol [2] (including the PIM-SSM variant [3]) is the main subject of the present document since it is most widely used in preference to other multicast routing protocols today.

PIM-SM over satellite is assumed to mean that PIM-SM relating to a given multicast flow is present at a BSM ingress router and is processed by the BSM network in such a way that the PIM-SM relating to the same flow is also made available to attached downstream networks at one or more BSM egress routers. It is also possible that PIM traverses the BSM system but only IGMP is then available downstream from the BSM egress ST, but this configuration is considered a minor variation as far as S-PIM is concerned and is not considered further here.

The PIM over BSM scenarios are outlined in [i.2] and in more detail in clause 4 below.

As PIM-SM is considered to be a control plane protocol, the main subject addressed by the present document function is the "Multicast Control" function as outlined in [i.2], rather than multicast IP flows in the user plane.

2 References

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.
- Non-specific reference may be made only to a complete document or a part thereof and only in the following cases:
  - if it is accepted that it will be possible to use all future changes of the referenced document for the purposes of the referring document;
  - for informative references.

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

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

2.1 Normative references

The following referenced documents are indispensable for the application of the present document. For dated references, only the edition cited applies. For non-specific references, the latest edition of the referenced document (including any amendments) applies.

[1] ETSI TS 102 294: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM) services and architectures; IP interworking via satellite; Multicast functional architecture".

[2] IETF RFC 4601: "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)".

[3] IETF RFC 4607: "Source-Specific Multicast for IP".

[4] ETSI TS 102 293: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM) services and architectures; IP Interworking over satellite; Multicast group management; IGMP adaptation".

2.2 Informative references

The following referenced documents are not essential to the use of the present document but they assist the user with regard to a particular subject area. For non-specific references, the latest version of the referenced document (including any amendments) applies.

[i.2] ETSI TS 102 461: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM); Multicast Source Management".

3 Definitions and abbreviations

3.1 Definitions

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

architecture: abstract representation of a communications system

NOTE: Three complementary types of architecture are defined:

- Functional Architecture: the discrete functional elements of the system and the associated logical interfaces.
- Network Architecture: the discrete physical (network) elements of the system and the associated physical interfaces.
- Protocol Architecture: the protocol stacks involved in the operation of the system and the associated peering relationships.

control plane: plane that has a layered structure and performs the call control and connection control functions; it deals with the signalling necessary to set up, supervise and release calls and connections

egress ST: ST at which an IP multicast flow (of user data) exits the BSM network

flow (of IP packets): traffic associated with a given connection-oriented, or connectionless, packet sequence having the same 5-tuple of source address, destination address, Source Port, Destination Port, and Protocol type

forwarding: process of relaying a packet from source to destination through intermediate network segments and nodes

NOTE: The forwarding decision is based on information that is already available in the routing table. The decision on how to construct that routing table is the routing decision.

ingress ST: ST at which an IP multicast flow (of user data) enters the BSM network

IP multicast: IP networking protocol that allows members of a specific host group to receive copies of the same IP datagram, identified by a reserved multicast address as the IP destination address

IP multicast address: one of a range of IETF-defined addresses for multicast

NOTE: For IPv4 this corresponds to the range from 224.0.0.0 to 239.255.255.255.

IP host group: set of IP receivers for a given IP multicast group

Network Control Centre (NCC): equipment at OSI Layer 2 that controls the access of terminals to a satellite network, including element management and resource management functionality

user plane: plane that has a layered structure and provides user information transfer, along with associated controls (e.g. flow control, recovery from errors, etc.)

3.2 Abbreviations

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

ASM	Any-Source Multicast
BMAC	BSM Multicast Access Control
BMS	BSM Management System
BSM	Broadband Satellite Multimedia
BSM_GID	BSM Group IDentity
BSM_ID	BSM IDentity
BSR	BootStrap Router
CPN	Customer Premises Network
C-RP	Candidate-Rendezvous Point
EUG	End User Group
FDMA	Frequency Division Multiple Access
GID	BSM Group ID address
GW	GateWay
IETF	Internet Engineering Task Force
IGMP	Internet Group Message Protocol
INT	Internet/MAC Notification Table
IntServ	Integrated Services
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISP	Internet Service Provider
MAC	Medium Access Control
MACC	BSM Multicast Access Control Client
MACD	Multicast Access Control-D
MACS	BSM Multicast Access Control Server
MAM	BSM Multicast Address Management
MAMC	BSM Multicast Address Management Client
MAMS	BSM Multicast Address Management Server
MAR	Multicast Address Resolution
MBGP	Multicast Border Gateway Protocol
MCM	BSM Multicast Control Management
MCMC	BSM Multicast Control Management Client
MCMS	BSM Multicast Control Management Server
MER	Multicast Edge Router
MGID	Multicast Group Identification Number
MMT	Multicast Mapping Table
MSDP	Multicast Source Discovery Protocol
NCC	Network Control Centre
NMC	Network Management Centre
OBP	On-Board Processing
OMCS	On-demand multicast connection setup
OUI	Organizationally Unique Identifier
PID	Packet IDentifier

PIM	Protocol Independent Multicast
PIM-ASM	Protocol Independent Multicast - Any-Source Multicast
PIM-SM	Protocol Independent Multicast - Sparse Mode
PIM-SSM	Protocol Independent Multicast - Source Specific Multicast
QID	Queue IDentifier
QoS	Quality of Service
RP	Rendezvous Point
RPF	Reverse Path Forwarding
RSM-A	Regenerative Satellite Mesh - A
SD	Satellite Dependent
SDAF	Satellite Dependent Adaptation Functions
SI	Satellite Independent
SIAF	Satellite Independent Adaptation Functions
SI-SAP	Satellite Independent Service Access Point
S-PIM	BSM Adaptation of PIM
SSM	Source Specific Multicast
ST	Satellite Terminal
UT	User Terminal
VP	Virtual Port

4 PIM-SM Scenarios in BSM

The scenarios described here represent the main ways in which PIM-SM can be employed and configured across a BSM system, which together with its attached networks forms part of an overall IP network.

The two modes of operation of PIM-SM are Any-Source Multicast (ASM) and Specific-Source Multicast (SSM). For example, Join messages in ASM specify only the Group address, G, (not the source address, S) as symbolised by the couple (*, G), while Join messages in SSM specify (S, G). Join messages in ASM are forwarded up the tree towards the Rendezvous Point router (RP), while in SSM they are forwarded to the Source.

ASM is often the default mode, and this mode can transition to the SSM mode during the multicast session.

The relative positions of source and RP for ASM mode play a role in the message flows across the BSM as described in the following clauses.

4.1 Source and Rendezvous Point Scenarios

These scenarios describe ASM and SSM modes, and complement those described in [i.2].

4.1.1 Source and RP in Core Network

A typical scenario for BSM is where the BSM is used as an access network to the Internet and both the source and RP are located in the core network.

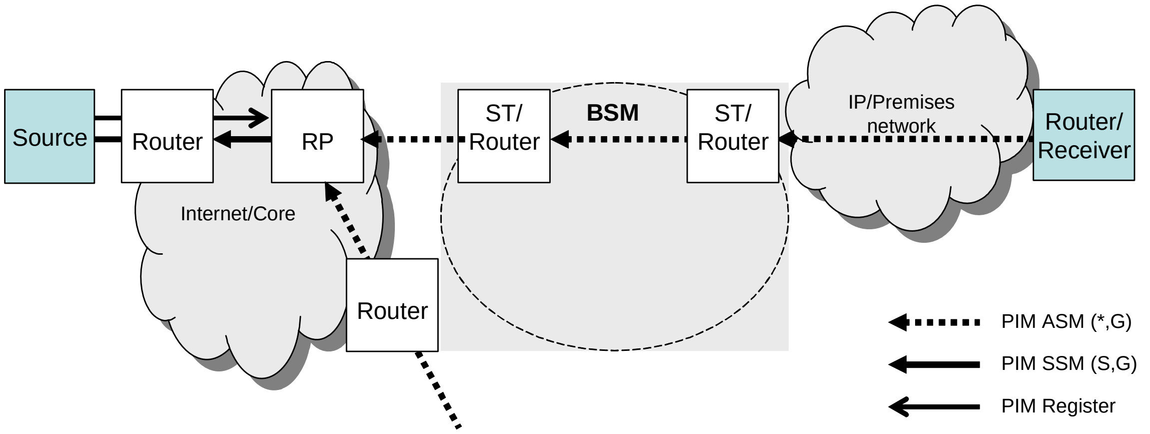

Figure 4.1: PIM-ASM mode, Source and RP in Core Network

The case of an RP located in the CPN with a source in the core is considered unrealistic, and is not discussed here.

4.1.2 RP and Source in Premises Network

In some cases a source may be located in, or attached to, the CPN. Any PIM messages from remote terrestrial network routers would have to be forwarded over the satellite to the RP, leading to increased traffic over the BSM compared to an RP located in the core network. The option for an RP located in the CPN may therefore be disabled for certain sources or network configurations.

This scenario is, however, well suited to multicast trees solely within the satellite network.

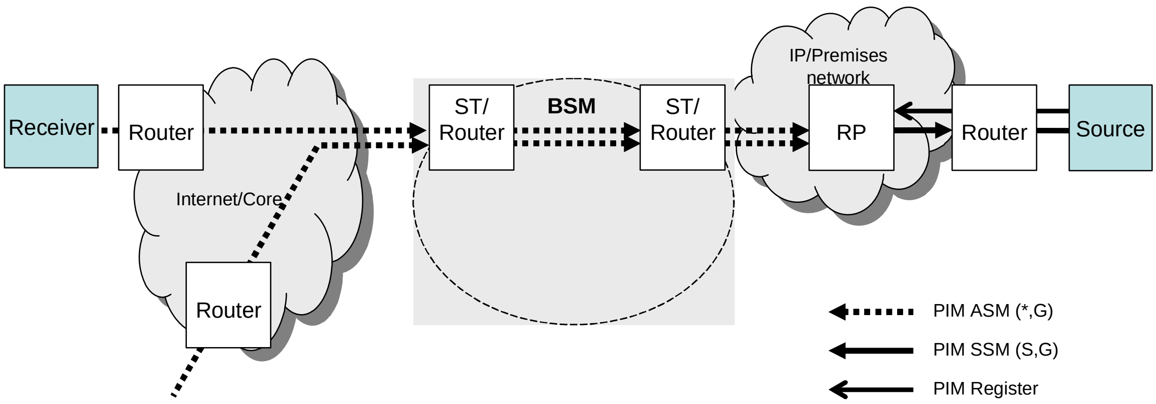

Figure 4.2: PIM-ASM mode, RP and Source in Premises Network

4.1.3 Source in Premises Network, RP in Core Network

As indicated in the previous clause, this scenario avoids PIM messages from remote terrestrial network routers being forwarded over the satellite. However, multicast data flows are forwarded by the source to the RP over the satellite irrespectively of whether any receivers are interested in the source data. This may result in unnecessary capacity being used.

PIM-SSM messages (instead of PIM-ASM) traverse the BSM from the RP to the source, or alternatively unicast register-encapsulated multicast data is forwarded from the source to the RP.

Also it is not an optimal scenario for multicast confined to within the BSM to have an RP outside of the BSM, due to the additional paths incurred.

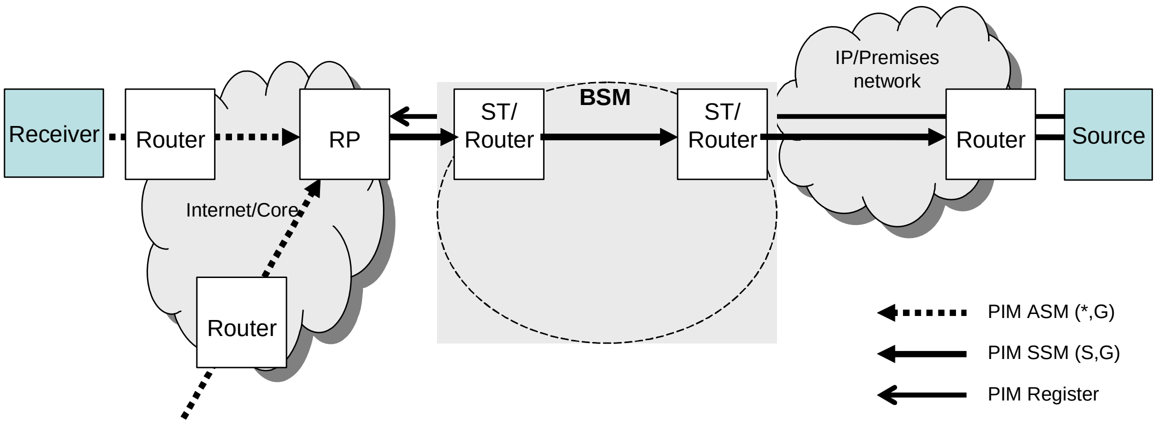

Figure 4.3: PIM-ASM mode, Source in Premises Network, RP in Core Network

4.1.4 PIM-SSM, Source in Core Network (or in CPN)

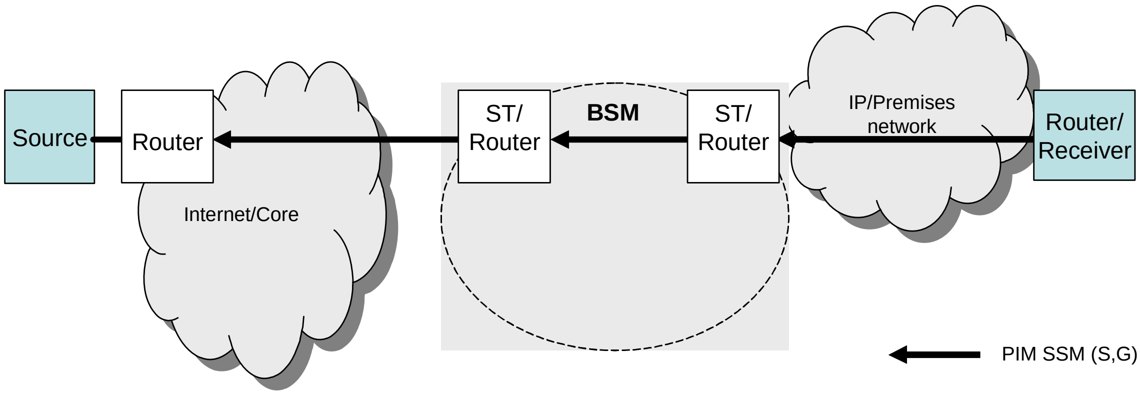

Figure 4.4: PIM-SSM, Source in Core Network

For SSM mode, only PIM-SSM messages traverse the BSM.

As far as the BSM is concerned, the case of a source located in the CPN is identical to this case in terms of PIM-SSM message processing.

4.1.5 Conclusion

In the above scenarios, all BSM ST's are considered identical as is the case in a mesh network. In this case the BSM can be considered symmetrical for PIM purposes, and the above scenarios can be reversed with respect to the BSM. Hence the PIM messages can traverse the BSM identically in either direction.

In the case of a star BSM network including a hub station, the same considerations apply as above providing that ST's and Hub Station have the same PIM functionality.

Both ASM and SSM mode messages should be similarly processed by the BSM without significant impact, as the main difference between these messages is the explicit specification of the source address for forwarding Joins etc. instead of the implicit RP address.

4.2 PIM Processing Options over the BSM

Scenarios considered for PIM-SM processing within the BSM are:

1) The BSM ST's act as PIM routers to external and internal BSM system interfaces (i.e. Native PIM-over-Satellite). The use of PIM-SM over the satellite implies that the ST PIM routers react dynamically and transparently to requests from networks downstream from the BSM, and to other types of PIM messages from both upstream and downstream networks. This also implies that connectionless IP routes are permanently available across the satellite system, or can be set up within a short reaction time.

2) The BSM ST's act as PIM routers to external interfaces, and the PIM-SM protocol would appear externally to transit the BSM system "transparently". Internally to the BSM an adaptation of PIM i.e. Adapted PIM-over-Satellite (S-PIM), or an alternative protocol could be used.

3) No PIM routing is implemented in the BSM ST's, but PIM snooping or proxying may be implemented to influence layer 2 switching, etc.

Since PIM is applied across the BSM network, the scenarios below are all considered to be of "Pull" type, in the terminology of [i.2], and of either Star or Mesh link topology between ST's.

These scenarios are discussed further below.

4.2.1 Scenario 1: Native PIM-over-Satellite

In this scenario, standard PIM messages are assumed to be forwarded in their native form by STs as routers within BSM networks. This clearly involves no adaptation of PIM in the sense of modification of the protocol compared to the standard.

The way in which the BSM acts upon PIM messages may depend on the SISAP and lower layer protocols and examples have been described in [i.2].

However, independently of such interactions, configuration of PIM within the bounds of allowed parameters may be advantageous.

The impact of PIM messages is considered below.

The general categories of messages between PIM-SM routers are:

- Hello
- Join
- Prune
- PIM Assert
- PIM Register/Register Stop
- BSR messages

The PIM-SM protocol is relatively garrulous and requires overhead for regular message passing to maintain PIM router state. This overhead may have unwanted impacts on a satellite system.

The Hello message is worthy of particular note, since it is sent when a multicast router starts up, and periodically thereafter (at 30 s intervals by default), on each PIM-enabled interface. The Hello message allows a router to learn about the neighbouring PIM routers on each interface, and perform other functions such as negotiating options. A router must record the Hello information received from each PIM neighbour. The Hello message is sent to the 'ALL-PIM-ROUTERS' multicast group address.

The Hello message is therefore likely to generate most PIM signalling traffic in a satellite system in which typically many egress routers are connected to each ingress router. Each message is at least 4 bytes long.

Also Join messages are typically sent regularly to keep alive the state. Each message is at least 7 bytes long.

Configuration of PIM-SM timers can be done to optimise performance over satellites and this is discussed in annex A.

4.2.1.1 BSM Link Architecture Impacts

PIM requires bi-directional multicast forwarding, in that multicast packets must pass in both directions over satellite links, because not only multicast user plane flows but also PIM control plane flows rely on multicast. This implies that satellite L2 links in both directions (even if not inherently multicast) must be configured to support L3 IP multicast.

Therefore for PIM to function correctly, the BSM must be fully meshed between all ST's and Hubs, i.e. bidirectional L3 paths available between each BSM ST/router to all other neighbouring ST/routers.

4.2.1.1.1 Star Architecture

The Hub has multicast L2 links to STs in the forward direction, but the STs do not generally have direct multicast links in the return direction to the Hub and all other STs. The return links can be made into multicast links effectively by rebroadcast of all PIM multicast flows received by the Hub on its satellite interface onto the forward link(s) (as well as onto other interfaces). This creates double-hop multicast links from STs.

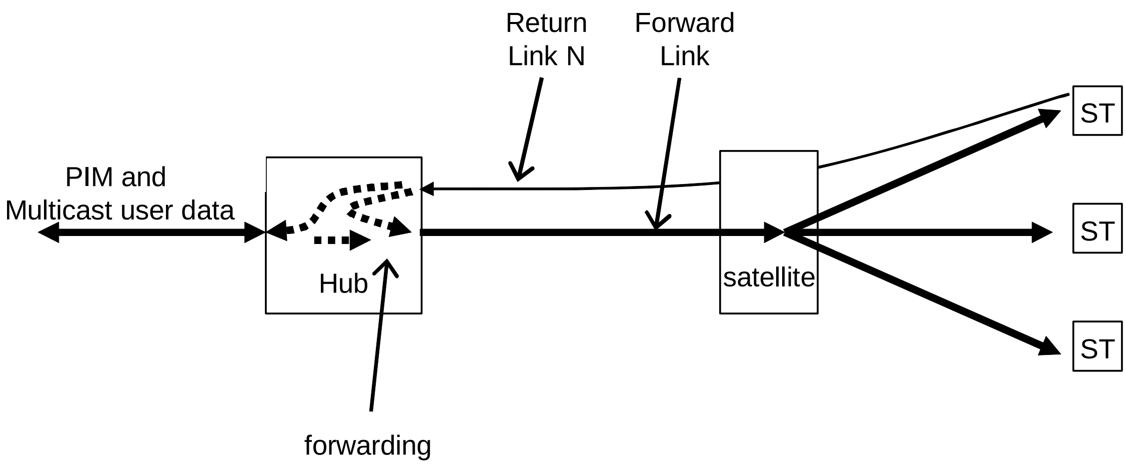

Figure 4.5: Forwarding in Hub for bidirectional PIM multicast

This architecture would give rise to multiplication of PIM messages over the satellite, but this is consistent with PIM.

4.2.1.1.2 Hybrid Star-Mesh (Transparent)

A way of avoiding such double hop links could be to use a hybrid satellite architecture using multicast satellite return links in several ways e.g.:

1) One dedicated L1 or L2 multicast satellite return link is available for all ST's. This could be accessed using ALOHA, for example, in a non-guaranteed way, which would be acceptable for PIM messages. This channel could be accessed on the downlink from the satellite by all STs (as well as the Hub).

2) A dedicated L1 or L2 multicast satellite return link is available per ST. This would have the disadvantage that each of the other STs would have to listen to all of these channels or links, of which there could be many. The number of links needed would also be larger.

These solutions would create an efficient mesh network in the return direction in a single hop. They could be restricted to multicast -only traffic, and PIM flows in particular.

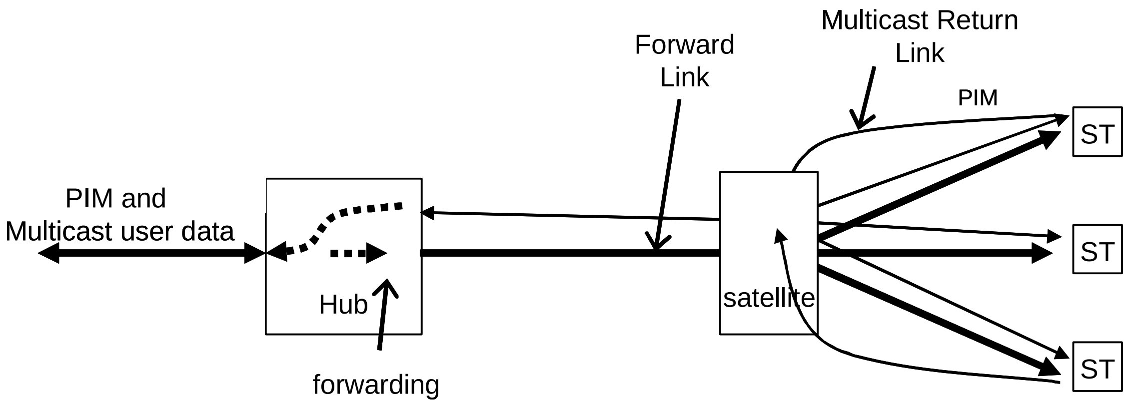

Figure 4.6: multicast mesh return link in hybrid "Star-Mesh" architecture for bidirectional PIM multicast (with single multicast return link)

The architecture of Figure 4.6 is only suitable when there are no multicast sources at the ST side. However this architecture is rarely, if ever, implemented. One difficulty is that more than one type of modulation and coding is needed on each link.

4.2.1.1.3 Hybrid Star Mesh (Regenerative Satellite)

With regenerative satellites, on-board demodulation and regeneration allows different uplinks and downlinks types to be flexibility interconnected, and single uplink and downlink formats to be used, resulting in the same connectivity as in Figure 4.6.

4.2.1.1.4 Mesh Architecture

In the mesh architecture, more than one ST generally act as ingress nodes. Each ST can forward multicast data (via the U-plane) to other STs directly over the BSM in a single hop without passing through the Hub.

Although the mesh architecture implies that all STs are directly connected to other BSM nodes, if this is not true for all nodes in the BSM (e.g. in the case of a very large number of STs) a hybrid architecture can be considered where lack of direct connections is replaced by indirect (double-hop) mesh connections through one or more nominated STs as described in clause 4.2.1.1.1.

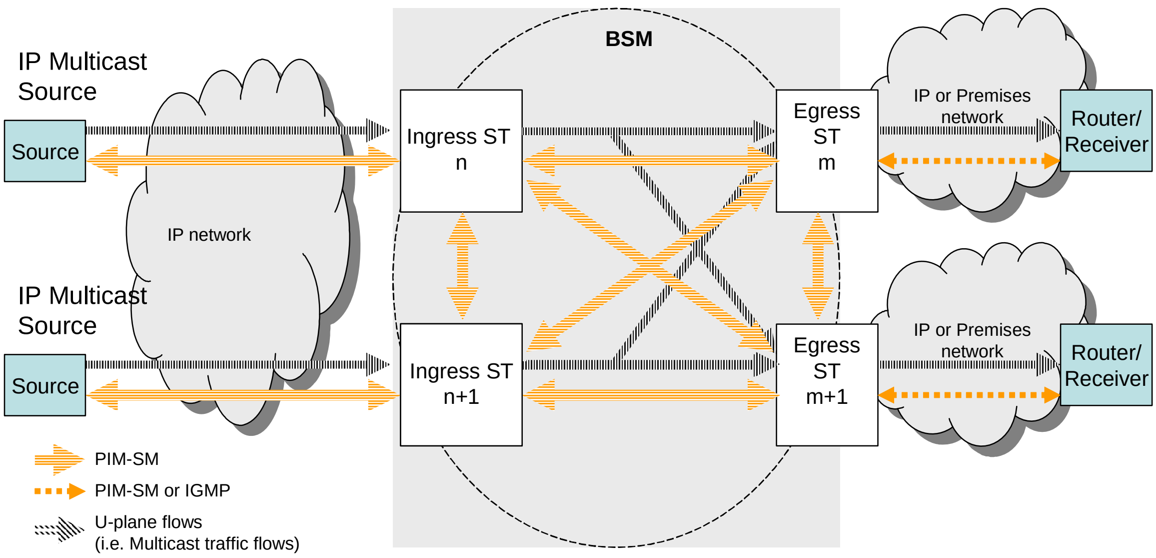

Figure 4.7: BSM native PIM Mesh Scenario showing message flows

NOTE: ST's can in general be both Ingress and Egress routers for different multicast Groups. They are shown as either Ingress or Egress for a particular Group in the above diagram for clarity.

4.2.1.2 Mesh Join Ambiguity and Assert Process

In the mesh scenario, the case of more than one ingress ST creates complications. For example two or more egress ST's may issue (*, G) or (S, G) Joins to different ingress ST's (because they have inconsistent MRIB entries regarding how to reach the RP or source S). Both paths on the RP tree will be set up, causing two copies of all the shared tree traffic to appear on the BSM.

These events have in general been foreseen in the PIM-SM standard, for the case of LANs.

PIM does not prevent such duplicate joins from occurring; instead, when duplicate data packets appear on the BSM from different routers, these routers should notice this and then elect a single forwarder. This election is performed using PIM Assert messages, which resolve the problem in favour of the upstream router that has (S, G) state; or, if neither has or both routers have (S, G) state, then the problem is resolved in favour of the router with the best metric to the RP for RP trees, or the best metric to the source to source-specific trees.

There must therefore be multicast mesh links between ingress ST's to communicate Assert messages.

The assert mechanism is only used when two STs can both forward traffic from the same remote source. The assert mechanism can resolve configuration errors or be used as a method to provide robustness to ingress ST failure.

4.2.2 Scenario 2: Adapted PIM-over-Satellite (S-PIM)

For native PIM signalling over satellite over mesh architectures as described in clause 4.1, there are potential difficulties especially from the egress STs towards the ingress: Joins and Prunes are sent on upstream interfaces to the "'ALL-PIM-ROUTERS " multicast address. Other PIM messages such as Hello are also sent on all interfaces to this address. Also PIM Assert messages are distributed between ingress ST's. This would lead to unnecessary multiplication of redundant messages, for example to STs which are not ingress routers (see Figure 4.7).

Star satellite architectures do not suffer so much (in terms of redundant signalling messages) as mesh architectures due to their single signalling route between STs and the Hub station; hence PIM adaptation would not be as attractive in this case.

In the mesh architecture therefore, for reasons of scalability, and as the number of unicast unidirectional mesh connections needed for full interconnectivity increases as n(n-1) for n STs, it was recommended in [i.2] that all C-plane messages in this scenario be adapted to a PIM star (or client-server) architecture as shown in Figure 4.8. The introduction of server for adapted-PIM messages could reduce the total number of messages over the satellite (to about 2n or fewer instead of n[n+1]) and may lead to other benefits in setting up satellite links for multicast. The disadvantage of introducing double hop links between ST's per PIM message would not be a significant disadvantage for delay in messages of this type.

In this Scenario therefore, the BSM ST's are configured to act as PIM routers to external interfaces, and to the outside world the PIM-SM protocol acts as though it transits the BSM system "transparently". Internally to the BSM an adaptation of PIM or an alternative protocol would be used.

Hence, for example, when an ST issues a Join, this must be sent to a BSM PIM server (as part of a more generalised Multicast Control Management Server), which may then decide to issue a Join to only one of the Ingress STs. This involves an adaptation of PIM-SM as described in clause 5.

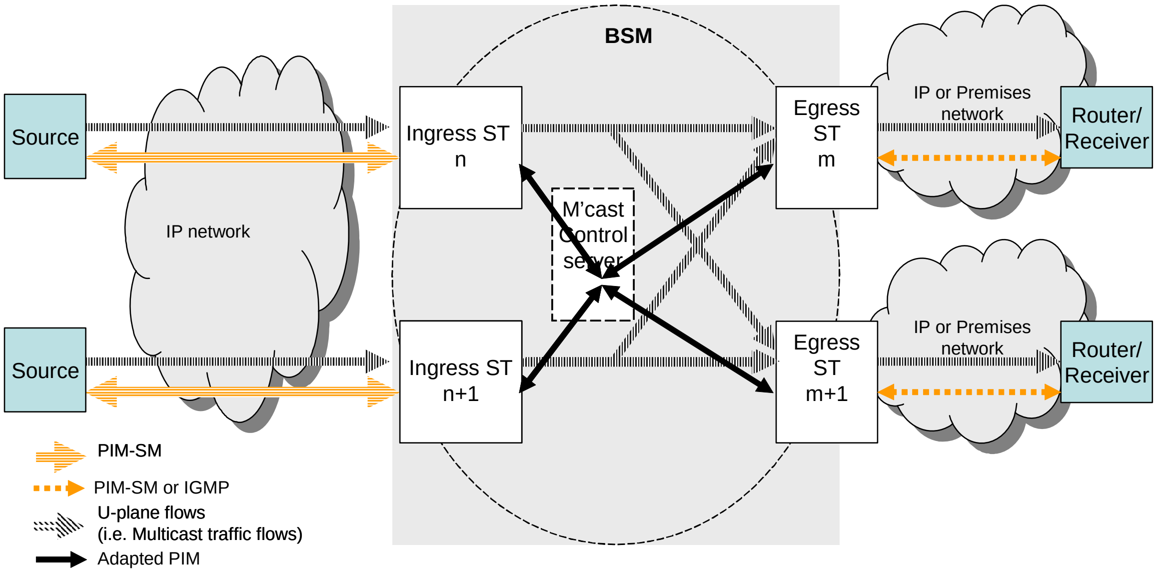

Figure 4.8: BSM Mesh Scenario - adapted PIM with BSM PIM server

All PIM messages are intended to be multicast, but if instead STs send these messages to the PIM server, the server can then multicast (or unicast) them to the appropriate groups (or a single ST) from one site instead of having to set up many trees.

Although in this scenario some small complexity may be added to the ST's due to protocol interworking, there are corresponding potential advantages in performance.

4.2.3 Scenario 3: PIM Snooping/Proxying

Where Layer 2 only is used to interconnect STs, (i.e. PIM-SM messages transit ST's transparently) the network floods IP multicast packets on all multicast router ports by default, even if there are no multicast receivers downstream. With PIM snooping or proxying enabled, the network can be used as a virtual switch to restrict multicast packets for each IP multicast group to only those multicast router ports that have downstream receivers joined to that group. The switch learns which multicast router ports need to receive the multicast traffic within a specific tree by listening to the PIM hello messages, PIM join and prune messages, and bidirectional PIM designated forwarder-election messages.

Further details of PIM snooping or proxying solutions are beyond the scope of the present document.

5 BSM PIM Adaptation (S-PIM)

This clause specifies the adaptation of PIM for use within the BSM when a PIM server is employed (i.e. taking the Mesh Pull case as the functional baseline). It is applicable to the mesh satellite system architecture as described in the scenarios of clause 4.2, and in which all STs are assumed to be PIM-enabled.

5.1 S-PIM over BSM Functional Architecture

This clause defines firstly the BSM S-PIM Functional Architecture. The PIM-SM protocol configures forwarding of the multicast flows for multicast groups and sources (in the User Plane) and hence PIM is considered here to be a Control Plane process, which is the focus here.

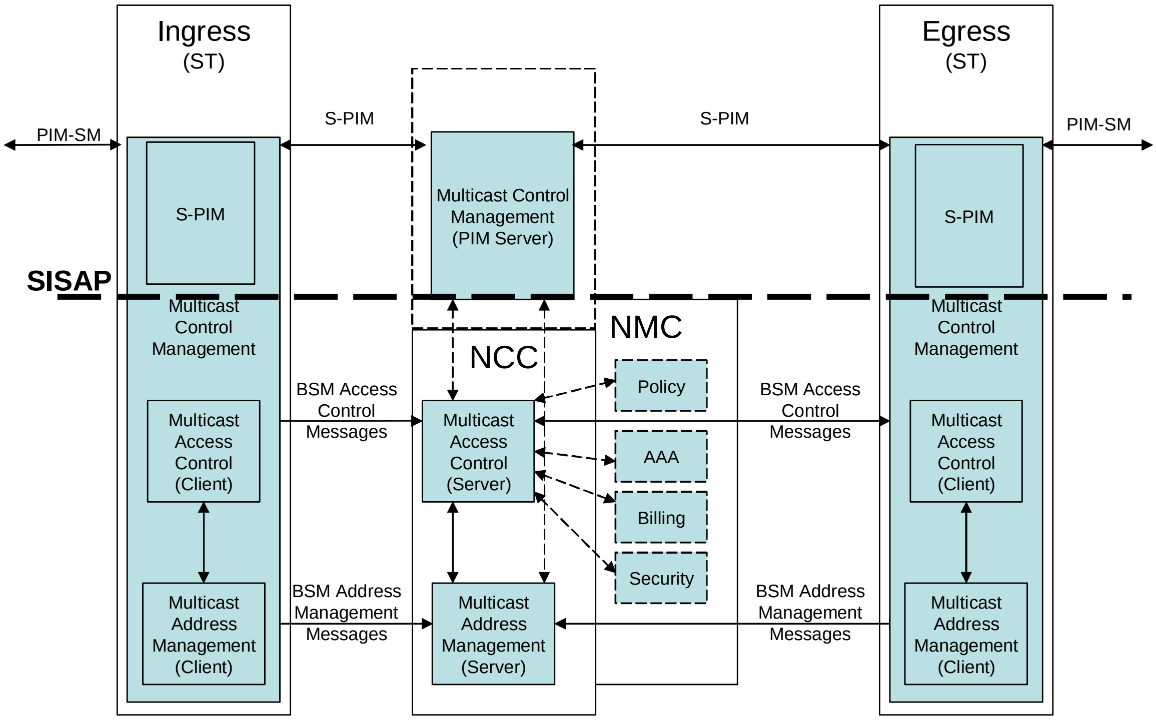

Figure 5.1: BSM Multicast Control Plane Architecture with internal Adapted PIM (S-PIM)

The PIM Server is considered part of a more generic BSM functional block (the Multicast Control Management function) and interacts with multicast control functions below the SISAP, which in this case are assumed to be located in the NCC and NMC.

5.2 Overall S-PIM Processing

Many options can be proposed for S-PIM processing depending on the generality of, or limitations on, multicast networking and sources.

Two cases are considered below for S-PIM processing:

1) S-PIM with normal egress-initiated set-up

2) S-PIM with ingress-initiated set-up

These cases are described further below.

In either case all S-PIM messages are passed between ST's and the PIM Server only. Firstly native (multicast) PIM messages issued by ST's within the BSM are converted into unicast messages directed to the PIM Server. The nature of the S-PIM messages depends largely on the method of passing these messages between the entities involved, namely the following options:

1) Native PIM packets from ST's are encapsulated into S-PIM packets

2) Native PIM packets from ST's are translated into S-PIM packets

There is little difference in processing between these schemes apart from the S-PIM message format. The aim is to retain the native PIM message format within S-PIM as much as possible by transparent encapsulation of such messages, for maximum compatibility with existing processes.

5.2.1 Case1: S-PIM with normal egress-initiated set-up

The most general case of PIM operation is where ST's act in the closest way to standard PIM routers and the PIM Server is virtually transparent to PIM messages; PIM messages received from STs are extracted from S-PIM messages and are multicast to all STs without modification (Figure 5.2).

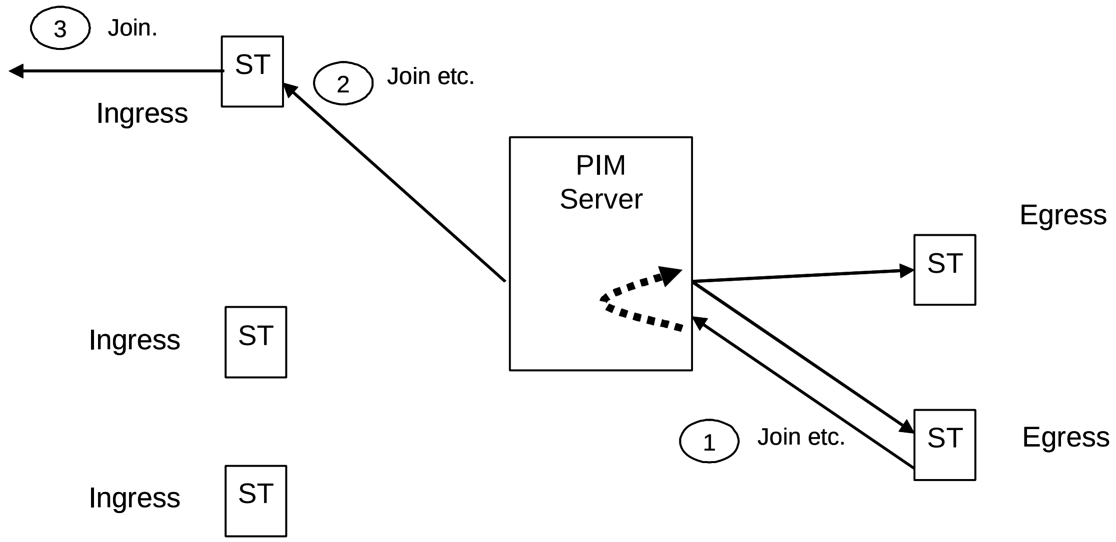

Figure 5.2: Message flow for general S-PIM case with multicast forwarding of PIM messages

As this case relates to connectionless multicast services, unless there is a mechanism for ensuring sufficient QoS, it is best-suited for applications with a high tolerance for latency and jitter, such as non-real time data transfer applications. Otherwise only a restricted set of Groups may allowed to conform with resource constraints

In principle all ST's can behave as both ingress and egress routers, but as in conventional PIM-SM, Joins are accepted for further upstream relaying only by a router which has been elected as the designated ingress router for that (S,G) or (*,G) (e.g. through the Assert process).

The resource requirements aspects of this case are for further study.

5.2.2 Case2: S-PIM with ingress-initiated set-up

A specific case of reverse PIM operation in which the ingress ST initiates multicast service set-up by negotiation with the PIM server.

PIM Join messages received by the PIM Server from egress STs are extracted from S-PIM messages in the normal way, but these requests are stored in the PIM Server and not passed further.

Independently one or more ingress STs are pre-configured by the NMC or PIM Server to receive only certain Groups (which have known QoS requirements). These STs issue Joins upstream to potentially receive any or all of these multicast flows. If an Ingress ST then receives one of these multicast flows it issues a Multicast Set-up Request (MSR) to the PIM Server.

The PIM Server checks to see if a Join from an egress ST has been stored for this Group and if so proceeds to set up the forwarding service to specific STs. (Figure 5.3), and sends the MSI message to the ingress ST.

This case is better suited to multicast applications with stringent quality of service requirements generally characteristic of real-time multimedia applications. It requires the PIM Server and STs to act in a more BSM-specific manner, and imposes some multicast configuration restrictions.

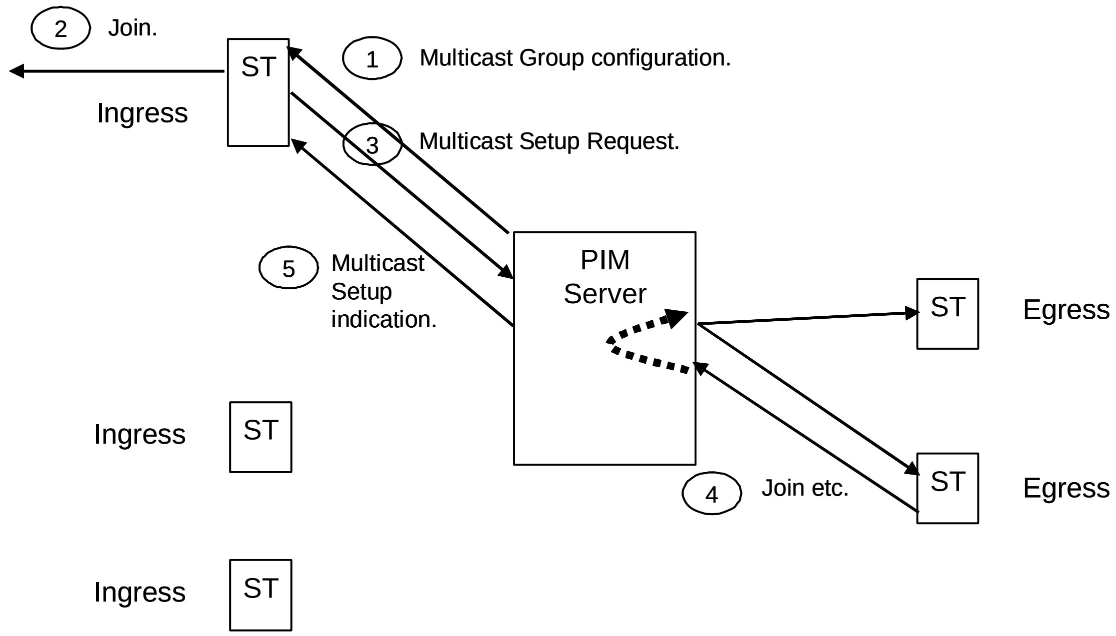

Figure 5.3: Message flow for special S-PIM case

In the above cases of S-PIM operation therefore, the BSM is ideally configured such that there are permanent links from ST's to the PIM Server to carry S-PIM signalling traffic. Otherwise on-demand links would have to be set up one by one as needed with inevitable signalling delay.

From the PIM Server to all STs, the BSM should also be configured with a permanent multicast link to carry S-PIM signalling traffic for both cases above. As a second option (for this case), permanent or on-demand unidirectional (unicast) links from the PIM server (to certain STs) could be sufficient.

Otherwise this architecture requires configuration of control plane routing and/or modification of the PIM protocols as defined in [i.2].

In the following Case 2 above will be described in more detail, since Case 1 follows conventional PIM operation.

5.3 S-PIM Message Processing and Forwarding

N.B. The differences between PIM and S-PIM concern only the internal BSM PIM messaging and processing. The following descriptions concern only the changes to standard PIM. Otherwise internal PIM messages and processes should be assumed to be unchanged.

Ingress STs are preconfigured with multicast groups of interest (Class D group IDs) by the NMC. These Ingress STs issue PIM Joins to the external network to try to join the groups in question. If an ingress ST has joined any of these groups, it monitors the relevant Class D group ID of any multicast packet received on its terrestrial interface and, if the group is already configured for forwarding, forwards the packet. If the Ingress ST is not yet configured for forwarding, the ST issues a MSR to the PIM Server.

When the Multicast Setup Request from an ingress ST is received, the PIM Server processes the request as follows:

1) verifies if the ingress ST is permitted to initiate a MSR;

2) verifies if it has received a Join from an ST in the multicast group for which the connection setup is being requested. If there is one or more receiving STs in the group who have joined, the PIM Server proceeds with the processing of the setup request and issues an MS Indication (MSI) message to the Ingress ST. If there is no active egress ST, the PIM Server sends a "Setup Release" message to the ingress ST;

3) verifies with the NCC/NMC that layer 2 parameters will allow the set-up, for example:

   - that system capacity is available based on the associated request parameters;

   - configuration of the satellite payload, if required.

4) if all these checks are completed successfully, the PIM Server grants the MSR and sends confirmation and setup information to the ingress ST.

While the MSR process is in progress in the PIM Server, the ingress ST discards multicast flows it receives for the group. When MSR is completed, the multicast service begins and the ingress ST forwards multicast datagrams received to the group.

5.3.1 RP and BSR Processes (PIM Server, etc.)

S-PIM multicast network administration requires that the BSR-set and the RP-set are on the same side of the BSM network as the ingress PIM ST. The PIM server stores the current RP information for each multicast session. The ingress PIM ST detects when the RP changes since all STs listen to BSR messages. The content of the BSR message indicates which RPs are alive. When an RP is down, it is removed from the BSR message. When an RP for a multicast session changes, the BSR sends out a new RP set information. When the ingress PIM ST detects a change in RP for a session, the ingress ST sends an RP Change Notification, in the form of an S-PIM management message, to the PIM server with the BSR message encapsulated in the management message. The PIM server transparently relays the encapsulated RP Change Notification as part of a management message and forwards it to the multicast group.

The PIM server includes the RP information in Join Acknowledgement/Announcement message that is sent to the multicast group when a Join is received. The Join Acknowledgement/Announcement message contains encapsulated BSR message if the RP for the group has changed. This allows PIM STs to update their RP information and keep it current.

When an egress PIM ST receives the announcement, it decapsulates the BSR message, checks and updates its RP information if necessary and forwards it on the ST user port(s) for which the RP's group prefix is valid.

The ingress ST maintains an "RP Change" state per group, which is flagged when an RP changes. This allows the ingress ST to process PIM Join/Prune messages and join or prune multicast delivery trees even though the RP in the PIM Join/Prune message received from an egress ST, via the PIM server, may not match the RP for the specified group in its route entry.

5.3.2 Other PIM Server functions

The PIM Server maintains several Timers, as for a normal PIM router:

- Join Timer
- Prune-Pending Timer

The PIM Server issues PIM messages depending on the status of these timers as for a standard PIM router. The PIM Server also resets these timers depending on reception of the S-PIM messages as described below.

The PIM Server interacts with entities below the SI-SAP as indicated in Figure 5.1 in order to set up and configure resources.

5.3.3 S-PIM Prune Processing

After successful MSR, the ingress ST restarts a Multicast Session Timer for every datagram that it transmits, to timeout a session. When the timer expires, the source ST assumes that there are no more datagrams to be transmitted and sends an Multicast Service Disconnect message to the PIM Server, which frees all resources that it allocated to the session and sends an acknowledgment to the ingress ST, which likewise frees all resources allocated to the session.

6 S-PIM Message Sequence Charts

6.1 S-PIM (Case 1)

Figure 6.1 shows the message flows for S-PIM for Case 1 as described in clause 5.2.1 above. In particular it shows Join processing when Join is accepted for the first time. Prune is similar except service data flows are potentially stopped.

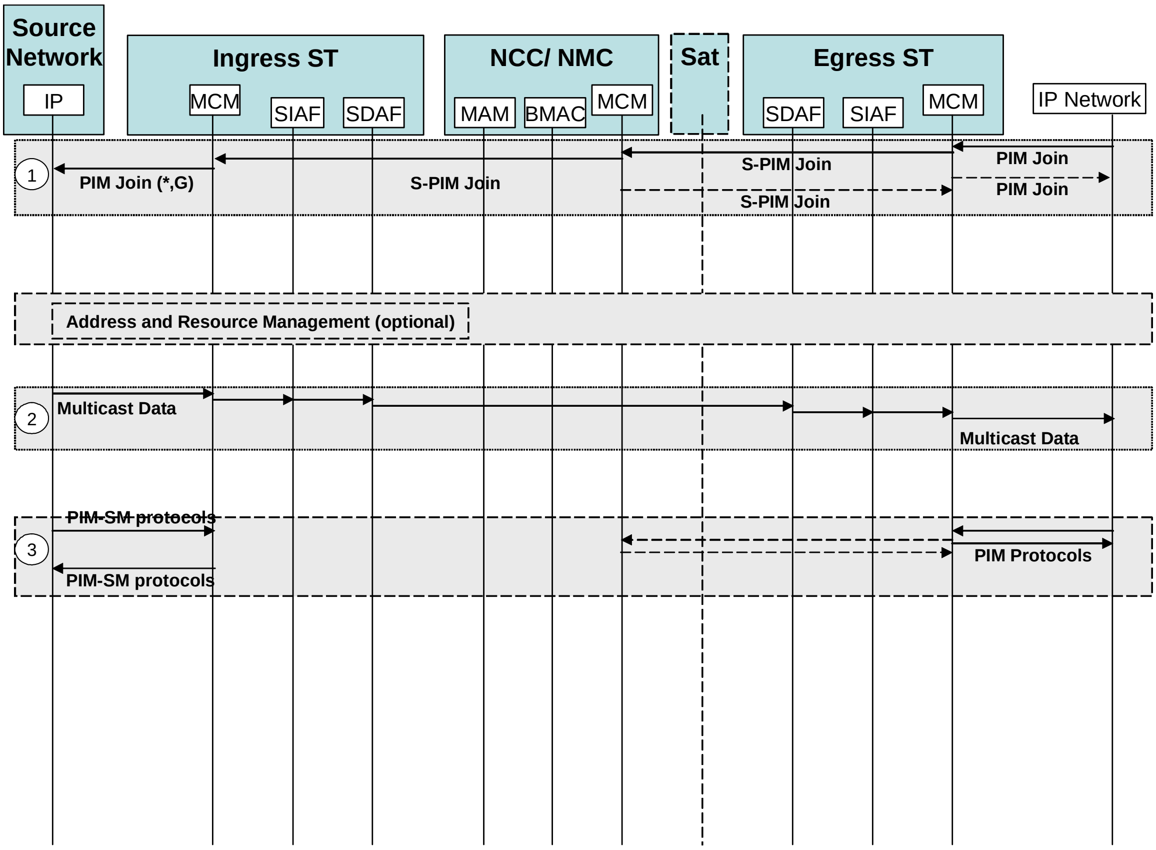

Figure 6.1: S-PIM Join Message Sequence (Case 1)

STEP	DESCRIPTION
1	In response to a PIM join (or an unsolicited IGMP "join" report) sent by a downstream network to an egress ST, the egress ST (MCM) adds a new group to its membership list (if necessary). It then issues a similar join message upstream; in this example this message is sent to the central PIM server, which acts upon the message. If this is a new multicast group for the BSM then the PIM Server requests authorisation from the BMAC server. If authorisation is granted, the PIM Server adds the group to its membership list and multicasts the Join to STs (or decides to which Ingress ST to pass the message).
2 & 3	Any subsequent signalling messages between the ST's are routed via the PIM server (e.g. Hello, Assert).

6.2 S-PIM (Case 2)

Figure 6.1 shows the message flows for S-PIM for Case 2 as described in clause 5.2.2.

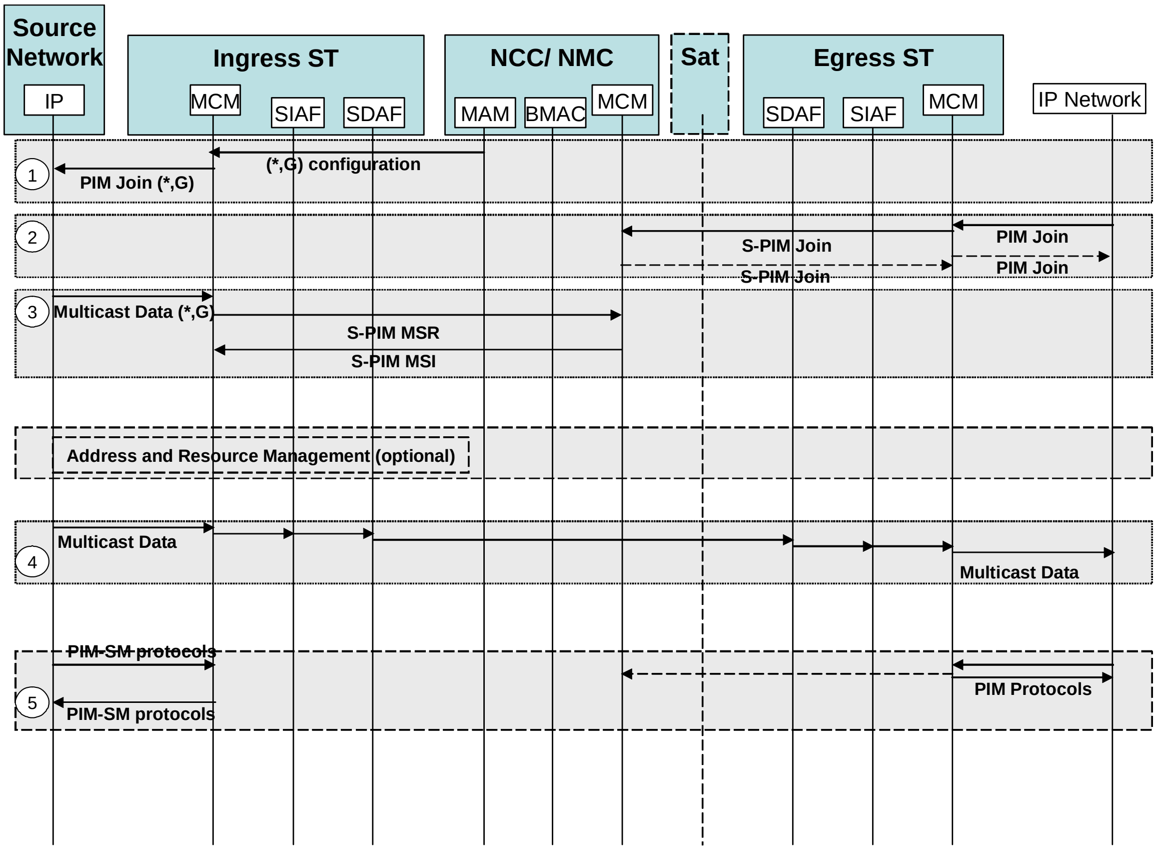

Figure 6.2: S-PIM Join Message Sequence (Case 2)

STEP	DESCRIPTION
1	Multicast Groups are pre-configured in the Ingress Routers by the PIM Server
2	In response to a PIM join (or an unsolicited IGMP "join" report) sent by a downstream network to an egress ST, the egress ST (MCM) adds a new group to its membership list (if necessary). It then issues a similar join message upstream; in this example this message is sent to the central PIM server, which acts upon the message. PIM Join messages received by the PIM Server from egress STs are extracted from S-PIM messages and the requests are stored in the PIM Server pending the setting up of a multicast tree if the group is already set up by an ingress ST, or pending a successful MSR.
3	If an ingress ST receives a new multicast flow it issues a Multicast Set-up Request (MSR) to the PIM Server. The PIM Server checks to see if a Join has been stored for this Group and if so proceeds to set up the forwarding service to specific STs.
4 & 5	Any subsequent signalling messages between the ST's are routed via the PIM server (e.g. Hello, Assert).

Annex A (informative): Configuration of PIM-SM over BSM

A.1 S-PIM Layer 2 Forwarding

Figure A.1 and Figure A.2 illustrate the flow of PIM messages using a BSM PIM server for two cases of link configuration in a mesh network:

1) for unicast links only;

2) for unicast return links and multicast forward links.

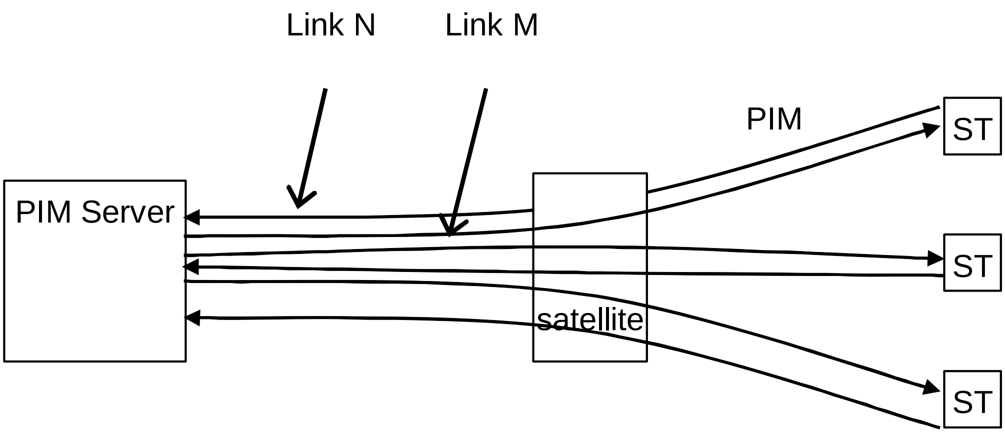

Figure A.1: Flows between STs and PIM Server with unicast links

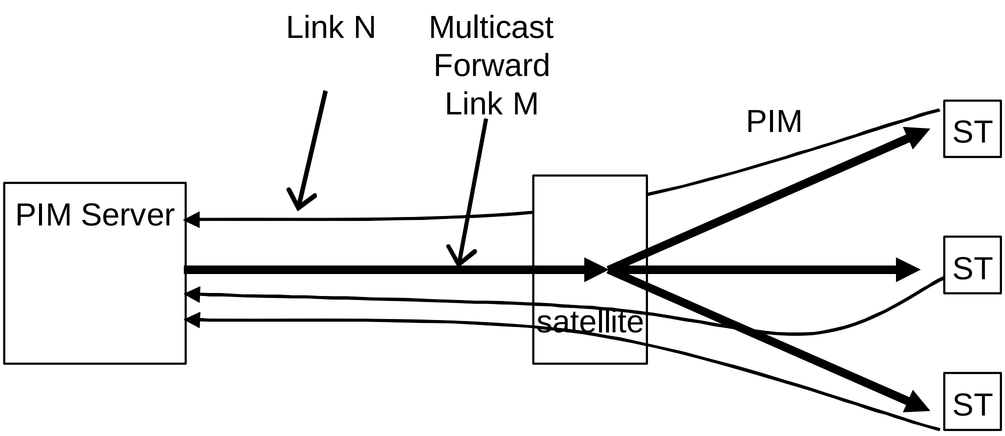

Figure A.2: Flows with multicast forward links

A.2 Reduction of Hello messages

One option carried in the Hello message is the Holdtime which is the amount of time a receiver of the message should keep a neighbour "reachable" i.e. open to accept other subsequent PIM messages . If the Holdtime is configured to '0xffff' for example, the receiver never times-out the neighbour, and avoids the need for periodic Hello messages.

Therefore in the BSM all egress routers may configure the holdtime on their satellite interface to reduce the system Hello message overhead, and in practice it may be prudent to avoid infinity holdtime and instead set the holdtime to a large value. In addition, the Hello_Period of the Hello Timer which determines the timing of Hello messages sent at each egress ST router's satellite interface can be set to a large value to reduce the Hello message frequency.

A.3 Reduction of periodic Join Messages

When there is more than one ST egress router attached to an ST ingress router and joined to the same (*,G) or (S,G) group, then the PIM standard recommends egress routers to suppress their own periodic Joins when they see a Join from another egress router, since Joins are multicast to the All_PIM_Routers group. This Join suppression becomes increasingly important as the number of egress routers in a group increases, and may well be significant in a satellite network.

This also implies that Joins can be multicast to all other egress routers in a satellite mesh network, or can be retransmitted by the ingress router in a satellite star network.

A.4 Reduction of Overrun of Multicast Transmission on "Prune" failure

A potential problem is waste of satellite resources because of latency in stopping multicast group transmissions when downstream multicast members or routers have left without reporting they have done so. This waste of resources could be significant particularly for wide bandwidth flows.

The lack of "prune" messaging may occur for several reasons, for example a downstream router failure or loss of a prune message.

If a Prune message from the last egress ST router to leave a group fails to arrive at a BSM ingress router, then the multicast stream may continue to be forwarded over the satellite.

A solution to such a problem is foreseen in the PIM-SM standard by means of the expiry_timer (ET) in any multicast router interface. When the interface is in a Join state and the ET expires, the interface reverts to the non-joined state even if no Prune has been received. The ET is reset to set to the maximum of its current value and the HoldTime defined by a Join message. Therefore to stay in the Joined state an interface typically needs to receive regular Joins to stop the timer expiring. The initial time of the ET is influenced as explained above by a parameter in the Join message called the Holdtime, which like the Hello message, can be configured to '0xffff' for which the interface never times-out. Clearly this latter option would be undesirable.

There is a trade-off between the multicast overrun allowed by the Expiry_Time and the frequency of Joins needed to maintain the Joined state, which are inversely proportional.

The frequency of regular Join messages issued by a downstream router in the Joined state is determined by the Join_Timer (JT) which by default is set to 60 s. The HoldTime is by default set to 3,5 * JT. Therefore the maximum overrun would normally be 3,5 minutes.

A.5 Prune Message Impacts

The removal of a Join state from an upstream router depends on one of two events:

1) Issue of a Prune from the last joined router on a downstream interface.

2) Expiry of the Join expiry_timer.

Annex B (informative): Bibliography

ETSI TS 102 357: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM); Common Air interface specification; Satellite Independent Service Access Point SI-SAP".

IETF RFC 1112: "Host extensions for IP multicasting".

IETF RFC 3569: "An Overview of Source-Specific Multicast (SSM)".

IETF RFC 4947: "Address Resolution Mechanisms for IP Datagrams over MPEG-2 Networks".

IETF RFC 5059: "Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)".

History

Document history		
V1.1.1	November 2009	Publication