ETSI TS 102 461 V1.1.1 (2007-01)

Technical Specification

Satellite Earth Stations and Systems (SES);
Broadband Satellite Multimedia (BSM);
Multicast Source Management

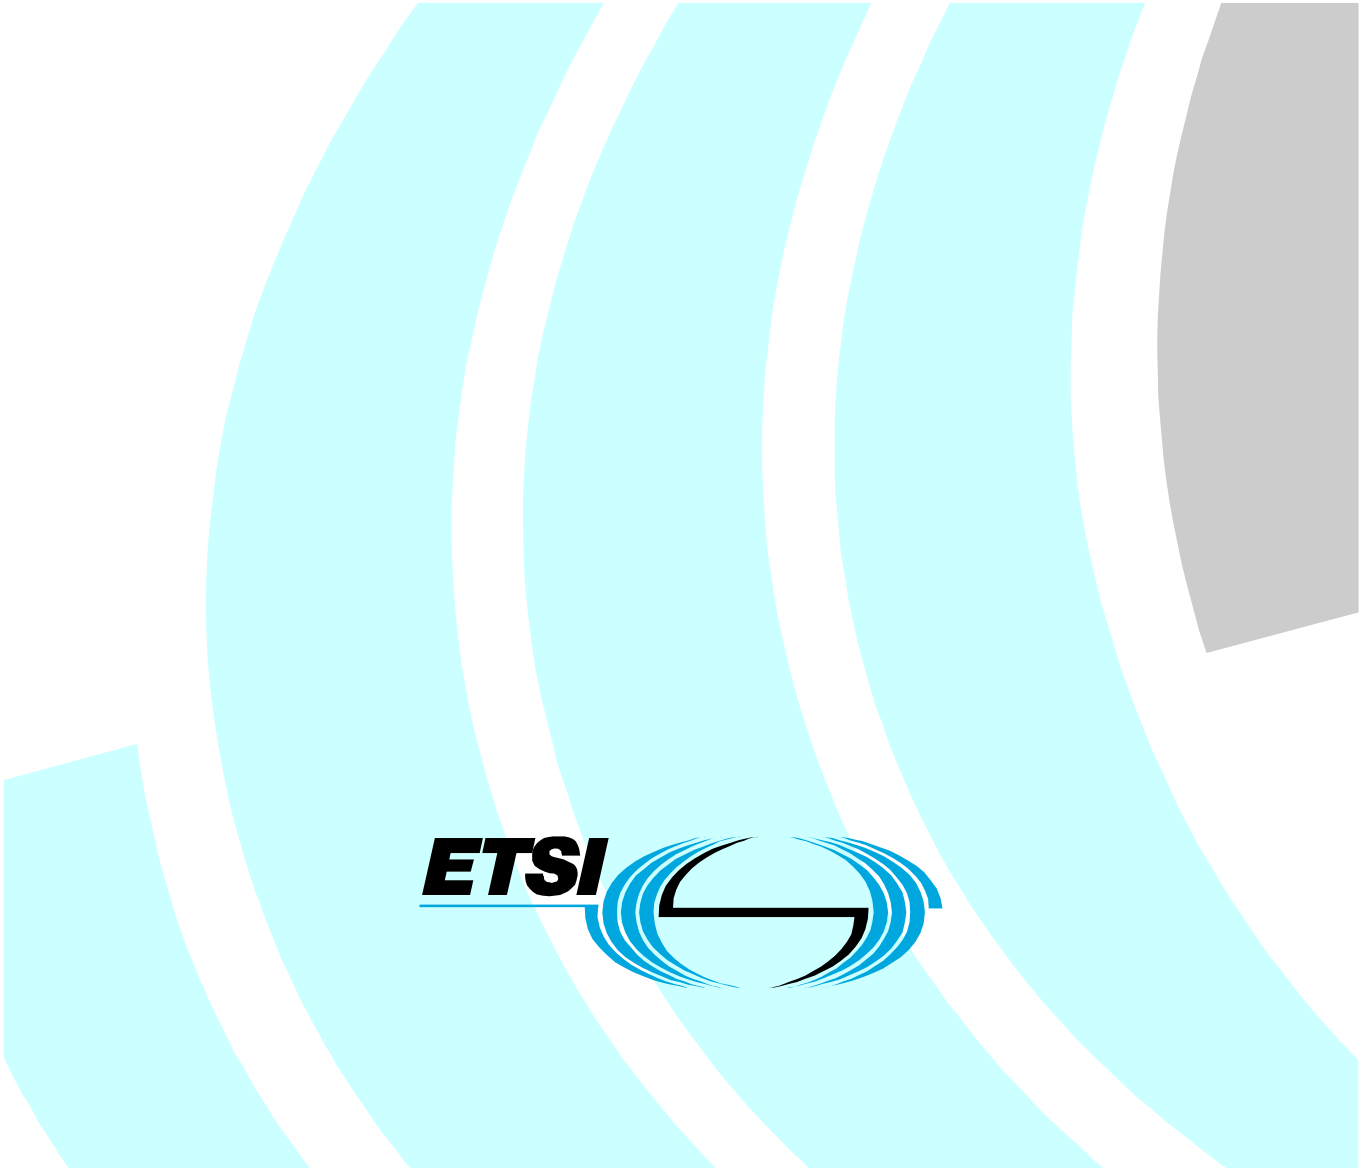

Reference
DTS/SES-00101

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

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

The present document defines how IP multicast sources, including local IP sources, may gain access to and enable their services be distributed across the BSM network. It considers where the multicast traffic is sourced, where it enters the BSM network, how it is transported across the network, and where it exits the network.

1 Scope

The present document defines the architectures and functions required for the interworking of IPv4 multicast protocols, including multicast sources, with the BSM.

The present document builds upon previous BSM documents referenced in clause 2, and notably TS 102 294 [1] on BSM Multicast Functional Architecture.

The present document firstly considers the BSM network scenarios for IP multicast interworking, which two main aspects:

1) the satellite network architecture;

2) the management of multicast sources and data forwarding, either statically or dynamically.

The BSM functional and protocol architectures are then derived for management of:

- IP multicast control messages (group management and routing protocols);
- Multicast access control (including resource management); and
- Multicast address resolution.

The present document then defines the detailed functional requirements and interactions of the above three functions with respect to the BSM lower layer interface, the SI-SAP. The Satellite-Dependent (SD) functions below this interface are system specific and are not treated here.

In the case of multicast routing protocols, the PIM-SM protocol [5] (including the PIM-SSM [I], [N], [O] variant) is taken as the basis for the present document since it is almost exclusively used in existing and proposed multicast routing applications today.

IPv6 protocols are not explicitly covered in the present document.

To make multicast an effective service over the BSM, multicast must take advantage of satellite's native multicast capabilities. Unlike Unicast, where destination IP and link layer addresses are specific to an end host, multicast employs a common IP "group" address for a given flow to all receivers, and therefore the BSM SI-SAP should also employ a corresponding common address, or GID (Group ID), for each multicast flow. The way in which these GIDs are controlled and employed is also defined in the present document.

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.

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

[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] 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".

[3] IETF RFC 1112: "Host extensions for IP multicasting".

[4] IETF RFC 3376: "Internet Group Management Protocol, Version 3".

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

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 Arhitecture:** 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

Differentiated services (DiffServ): services based on statistical (aggregate flows) guarantees and results in "soft" QoS

NOTE: Using packet markings (code points) and queuing policies it results in some traffic to be better treated or given priority over other (use more bandwidth, experience less loss etc.).

Egress ST: ST at which an IP multicast flow exits the BSM network

flow of IP packets: traffic associated with a given connection or connectionless stream 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 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

IP Multicast Distribution Tree: complete multicast distribution tree from IP source to all the IP receivers

Network Control Centre: 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:

AAA	Authentication, Authorization and Accounting
BMAC	BSM Multicast Access Control
BSM	Broadband Satellite Multimedia
BSM_GID	BSM Group IDentity
BSM_ID	BSM IDentity
BSR	BootStrap Router
COPS	Common Open Policy Service
CPN	Customer Premises Network
C-RP	Candidate-Rendezvous Point
DiffServ	Differentiated Services
DVB-RCS	Digital Video Broadcast - Return Channel for Satellite
EUG	End User Group
FDMA	Frequency Division Multiple Access
GID	BSM Group ID address
GRE	General Routing Encapsulation
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 Discoevery Protocol
MSP	Multicast Service Provider
NAT	Network Address Translation
NCC	Network Control Centre
NMC	Network Management Center
NOCC	Network Operations Control Center
NSIS	Next Steps In Signalling
OBP	On-Board Processing
OMCS	On-demand multicast connection setup
OUI	Organizationally Unique Identifier
PID	Packet IDentifier
PIM	Protocol Independent Multicast
PIM-SM	Protocol Independent Multicast - Sparse Mode
PIM-SM	Protocol Independent Multicast-Sparse Mode
PIM-SSM	Protocol Independent Multicast-Source Specific Multicast
QID	Queue IDentifier
QoS	Quality of Service
RCST	Return Channel Satellite Terminal

RGN	Replication Group Number
RP	Rendezvous Point
RPF	Reverse Path Forwarding
RSGW	Regenerative Satellite GateWay
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
SLA	Service Level Agreement
SMCS	Scheduled Multicast Connection Setup
SNMP	Simple Network Management Protocol
SSM	Source Specific Multicast
ST	Satellite Terminal
TDMA	Time Division Multiple Access
TTL	Time To Live
UT	User Terminal
VP	Virtual Port

4 BSM Network Scenarios

4.1 Summary

The BSM system is defined to be an IP compatible network and it is therefore required to interface and interwork seamlessly with external IP hosts, routers, and protocols.

The scenarios described below represent the main ways in which IP multicast can be employed and configured across a BSM system, which together with its attached networks forms part of an overall IP network. These instances of network configurations can be combined together in a real system implementation.

The scenarios are based on two main criteria of BSM network configuration and service control:

1) BSM network configuration - either star or mesh:

   - star topology - refers to a star arrangement of links between a central Hub station and remote STs through the satellite. The Hub acts as the sole BSM ingress router and distribution node for BSM multicast. The STs are all egress routers connected either directly to hosts or via premises networks;

   - mesh topology - refers to a mesh arrangement of links between STs where all STs can be interconnected directly through the satellite and each ST can act as a multicast distribution node to STs. STs can therefore be both ingress and egress routers.

NOTE: In both Star and Mesh cases, the satellite downlink coverage towards the egress STs can be split into several sub-groups by means of the Physical Layer configuration, for example by multiple satellite beam coverage channels and by other physical or MAC layer channels (e.g. TDMA, FDMA, PID). Multicast forwarding to each of these sub-groups can be made independently on each channel, which can be treated as one of a set of output ***interfaces*** for the ingress ST. Satellite OBP (On board processing) opens up further capabilities such as the ability to use one uplink channel for multicast from the Ingress ST, and replicate link layer streams in the satellite to selected downlink coverages.

2) Multicast service control - either push or pull

   - push - multicast services are configured by the BSM network operator, or similar centralized management entity, in terms of which groups are forwarded over the BSM on a quasi-static basis. The manager may not always know in advance what kind of resources (bandwidth, delay, jitter, etc.) will be required for a given multicast flow, but it has to configure BSM resources based, for example, on a service level agreement.

- pull - multicast services are requested and initiated dynamically (i.e. on demand) by each receiver host issuing a "join" to an IP multicast group, and therefore by relay though each egress ST, to the Ingress ST using IP multicast protocols. The conditions under which new group membership can be allowed and the associated multicast flows forwarded over the BSM are determined by BSM network policies.

There are therefore four main scenarios:

1) Star Push.

2) Star Pull.

3) Mesh Push.

4) Mesh Pull.

These four scenarios are not mutually exclusive and a given network can implement a combination of different scenarios for different groups.

There are some further options within these scenarios, for example depending on whether PIM-SM or IGMP is used either over the BSM or in the attached networks in each case.

The main features of the BSM network for each scenario are summarized in Table 4.1.

Table 4.1: BSM multicast network scenarios

Scenario	Multicast traffic Ingress Point	Multicast traffic Egress Point	BSM network IP multicast control	Ingress IP multicast control	Egress IP multicast control	BSM Access Control	BSM Address Management
STAR PUSH	Hub	ST	None	None/ IGMP/PIM	None/ IGMP/PIM	Static	Static/ Dynamic
STAR PULL	Hub	ST	IGMP/PIM	IGMP/PIM	IGMP/PIM	Dynamic	Dynamic
MESH PUSH	ST	ST	None	None/ IGMP/PIM	None/ IGMP/PIM	Static	Static/ Dynamic
MESH PULL	ST	ST	IGMP/PIM	IGMP/PIM	IGMP/PIM	Dynamic	Dynamic

A BSM distribution tree can form part of an IP multicast distribution tree, in the case of a public IP multicast group, or it may represent the complete distribution tree in the case of an internal (private) IP multicast group. These two options are also discussed below.

In all the following scenarios, a combined NCC/NMC (network control centre/network management centre) is shown combining all centralized BSM lower layer control and management functions. This entity may or may not be closely associated with the Hub station in star scenarios. In addition in some scenarios there may be a centralized IP layer multicast entity, the Multicast Control Management Server (MCMS), which may be associated with the NCC. The MCMS is not shown on the following diagrams for clarity, but is discussed in clause 5.

Any references to an operator are to be understood as being enacted via the NCC/NMC.

4.2 Star-Push

In the Star-Push scenario of Figure 4.1, all IP multicast services are pre-selected by the operator and the selected multicast flows are forwarded by the Hub over the satellite to all or certain Egress STs, depending on service agreements, and on the number of satellite beam coverages and channels.

The BSM multicast distribution tree is quasi-static, and IP or BSM multicast control protocols are not used by STs towards the Hub to request new multicast groups.

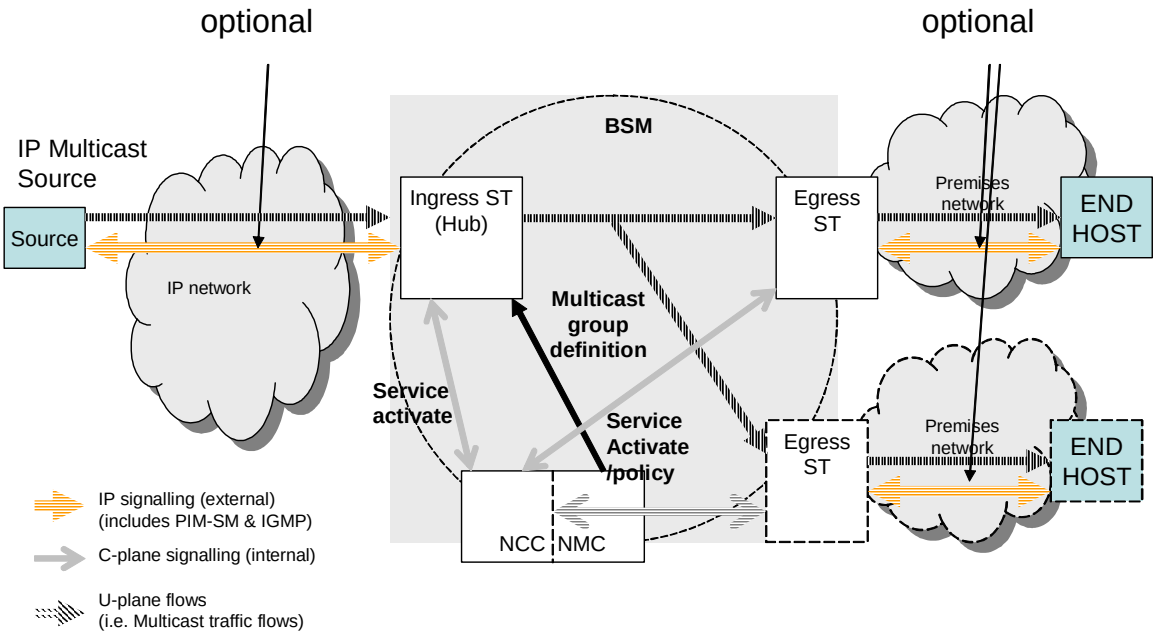

Figure 4.1: Star Push Scenario

4.2.1 Egress Nodes

All Egress nodes of the BSM are STs, and all those STs connected to a given Hub output interface can receive all the multicast flows on that interface.

Egress STs may process IGMP or PIM messages locally to control the forwarding of multicast groups. This implies that STs can have dynamic group membership but only to the groups to which the Hub is subscribed, and this membership has no impact on the traffic forwarded over the BSM.

The information needed by a given Egress ST to access the available multicast flows must be supplied via the NCC's multicast access control function. Such a function may be used to set the access policy at Egress STs to control the forwarding of multicast traffic on the premises network for potential reduction of traffic. This information can be sent, for example, either via common broadcast messages (e.g. in DVB-RCS the INT and MMT tables) or in individual signalling messages to each Egress ST. These same mechanisms can also be used to signal any changes to the available multicast groups and/or the policy and/or the associated distribution trees.

4.2.2 Ingress Node (Hub)

In this scenario the multicast ingress node to the BSM network is the Hub ST.

Group subscription of the Hub is done by configuration, indicated in the diagram by a connection to the combined NCC/NMC (network management centre). At the same time the NCC will use access control to set up the required lower layer multicast addresses and BSM distribution trees to STs.

IP multicast flows can then arrive at the Hub at any time once it is subscribed to a group. The Hub will forward the multicast data into assigned BSM distribution trees according to current policy.

Multicast traffic may arrive at the Hub/Ingress via a range of different paths and mechanisms. For example, it may be tunnelled from a remote source or arrive as native multicast packets via the attached network. The Hub may also process IGMP or PIM externally to initiate and maintain multicast group membership.

4.2.3 Multicast Sources from Remote STs

The remote ST variant of the Star-Push scenario concerns a multicast flow supplied from a remote ST via a unicast tunnel that transports the flow from the ST to the Hub. The Hub then extracts the flow from the unicast tunnel and forwards the resulting multicast data into the BSM distribution tree. As a result, the Hub is the source router for the multicast packets and hence the multicast aspects for this variant are identical to the baseline scenario.

In this variant a remote ST can send multicast/broadcast packets to a Gateway/Hub on an inbound link that subsequently forwards these natively on the outbound link, where they are finally received at the originating ST. This raises a need to carefully consider forwarding of these multicast packets, to avoid forwarding loops, (e.g. RFC 3077 [V] notes that a Receiver needs to be configured with appropriate filter rules to ensure it discards packets that originate from an attached network at the Egress and are later received over the outbound link). This filtering could utilize the L2 source address of the forwarding ST (if present) or the normal L3 source address (using a Reverse Path Forwarding (RPF), or other filtering methods.

4.3 Star-Pull

The Star-Pull scenario extends the Star-Push scenario to allow multicast control messages to transit the BSM from the Egress STs and thus allow dynamic group membership of the Hub to be initiated via the STs. The Hub can then re-configure the BSM distribution tree as needed, via the NCC. The NCC is responsible for making policy decisions on whether and under what conditions new group membership is allowed. It must also decide on resources to be allowed and the service level (QoS) to be provided.

These BSM multicast control messages can be IGMP or PIM, or an adapted version of these protocols, or a proprietary satellite dependent control protocol. In all cases, these control messages are internal to the BSM network: since they are internal they do not have to conform to any standard, as long as the ingress and egress STs can interwork with standard IP protocols externally.

In this scenario individual STs can request membership of groups, though the broadcast nature of the satellite forward link and the allocation of lower layer multicast addresses (i.e. BSM, GIDs, etc.) to IP groups implies, as before, that all STs connected to a given Hub output interface can receive all the multicast flows on that interface.

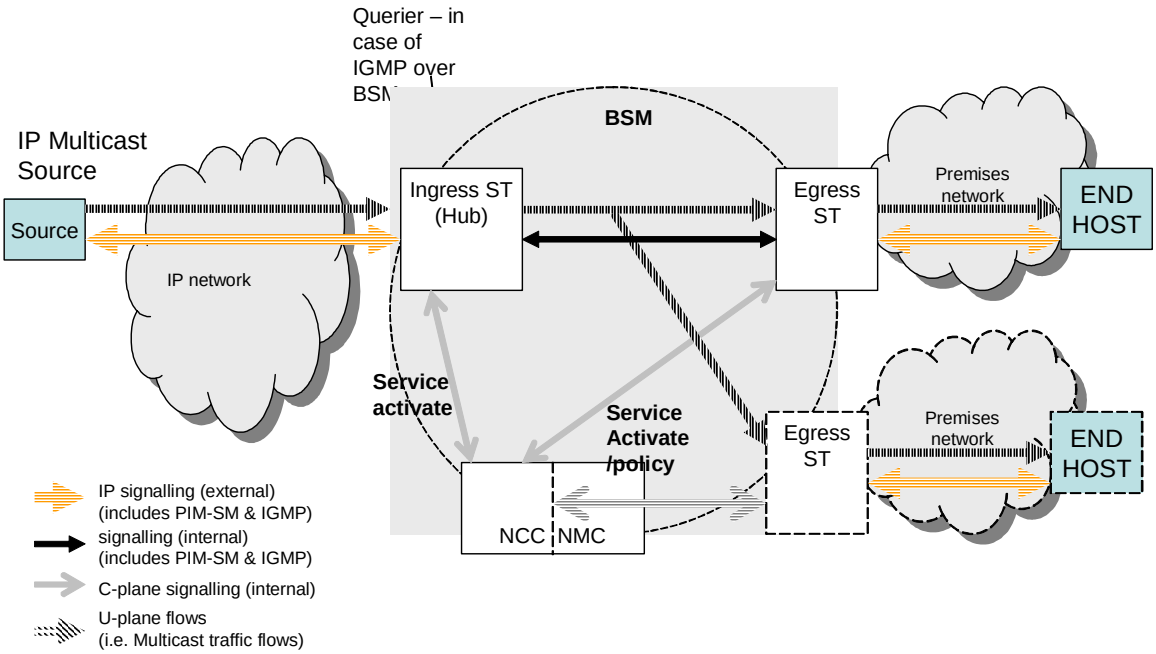

Figure 4.2: Star Pull Scenario

4.3.1 Egress Nodes

IGMP or PIM is used in the Egress ST's premises network to determine the IP group membership in these down-stream networks.

Clearly group membership of the ST is determined by the attached hosts and CPN, as well as policy from the NCC. Any change of ST group membership is signalled to the Hub via the internal control protocol (e.g. IGMP, PIM, etc.).

4.3.2 Ingress Point (Hub)

The Hub's IP group membership is determined by the attached STs as well as policy from the NCC, and in the case of multiple beam coverages the Hub may maintain a separate set of groups for each output interface.

Whenever a change of aggregated group membership across all its output interfaces occurs, the Hub signals to the external network to join or leave that group. Subsequently it requests the NCC to activate or de-activate access control (if necessary) and BSM multicast addressing for forwarding over the BSM.

4.3.3 Double-Hop Star-Pull Variant

A variant of the Star-Pull scenario concerns a Star BSM Network that is used to provide "mesh" multicast services at the IP layer from a remote Ingress ST to other Egress STs. The Hub may also be a member of the multicast group.

This scenario is similar to the Mesh-Pull scenario in terms of IP interworking, but the use of a Star BSM network means that the underlying links are obliged to use double-hop transmission via the Hub to forward multicast data from the Ingress ST to all the Egress STs.

This scenario is not considered further in the present document.

4.4 Mesh-Push

This scenario is essentially the same as that of Star Push except that more than one ST act as ingress nodes (if there is only one ingress ST then the multicast scenario reverts to Star Push). Each ST can multicast directly over the BSM without having to use the Hub in a star network. All STs connected to a given Ingress ST output interface can receive all the multicast flows on that interface.

The NCC/NMC is assumed to be a centralized server but in the mesh topology this is no longer co-located with the Hub or with any particular Ingress ST. This has consequences for the BSM architecture: the U-plane traffic and the associated distribution trees have a different topology to the C-plane control signalling.

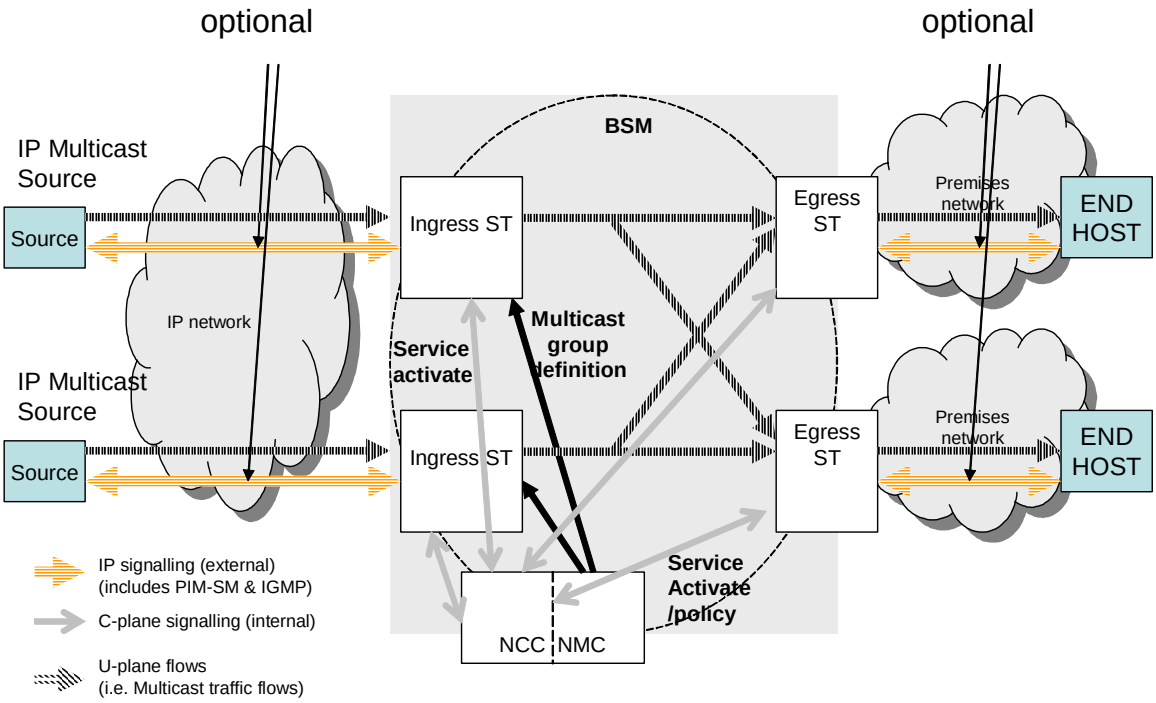

Figure 4.3: Mesh Push Scenario

4.4.1 BSM Network Multicast Forwarding

In this scenario there may be multiple ingress STs, which forward multicast flows arriving from different sources in the public internet, or from local domain sources. If these sources originate from the public internet, there is a risk they may use the same public group address and this may result in "collisions", or ambiguity, of group addresses at the Egress STs which receive these flows. Any ambiguity can be avoided in this case by configuration of the single relevant Ingress ST to receive the group and, in the case of PIM-SM, by specifying its upstream router (if there are several) or the RP of the group.

If needed, Egress STs can filter the flows they receive on their unicast IP source addresses to avoid collisions, but this may not provide sufficient resolution in all cases.

The most general method of separating multicast flows is to use SI-SAP address scoping to assign a different GID to each set of flows. This is further described in annex A.

4.5 Mesh-Pull

This scenario is similar to Star Pull except that more than one ST act as ingress nodes (if there is only one ingress ST then the multicast scenario reverts to Star Pull). Each ST can forward multicast data (via the U-plane) to other STs directly over the BSM (single hop) without passing through the Hub.

However for group membership signalling from the egress STs towards the ingress, there are potential problems:

a) To decide to which ingress ST to group membership messages are sent by any egress ST. IGMP join/leave reports are sent using either an "all routers" multicast address (224.0.0.22) or the group address concerned. It could be assumed that the BSM offers such a return point-multipoint bearer (e.g. using satellite OBP or double hop links). But generally the return link over a satellite offers unicast rather than point-multipoint capability, and either one, several, or all possible return bearers could be chosen in this case. A mechanism is then needed to choose one or more return links.

b) Group membership messages could be sent to all ingress STs (assuming a point-multipoint return link) but because of the nature of a satellite network this may result in possible unwanted outcomes, for example:

   - multiple Ingress ST memberships of the same group for forwarding to the same egress STs;

   - excessive signalling overhead to agree between ingress STs the optimum point-multipoint path and hence to choose the optimum ingress ST.

Therefore, all C-plane messages in this scenario are recommended to use the client-server architecture. Hence, for example, when an Egress ST (a client) issues a join request, the request must be sent to the Multicast Control Management Server (instead of the local router as is usual) which may then decide to issue a membership request to only one of the Ingress STs. This is discussed further in clause 5.2.1.4.

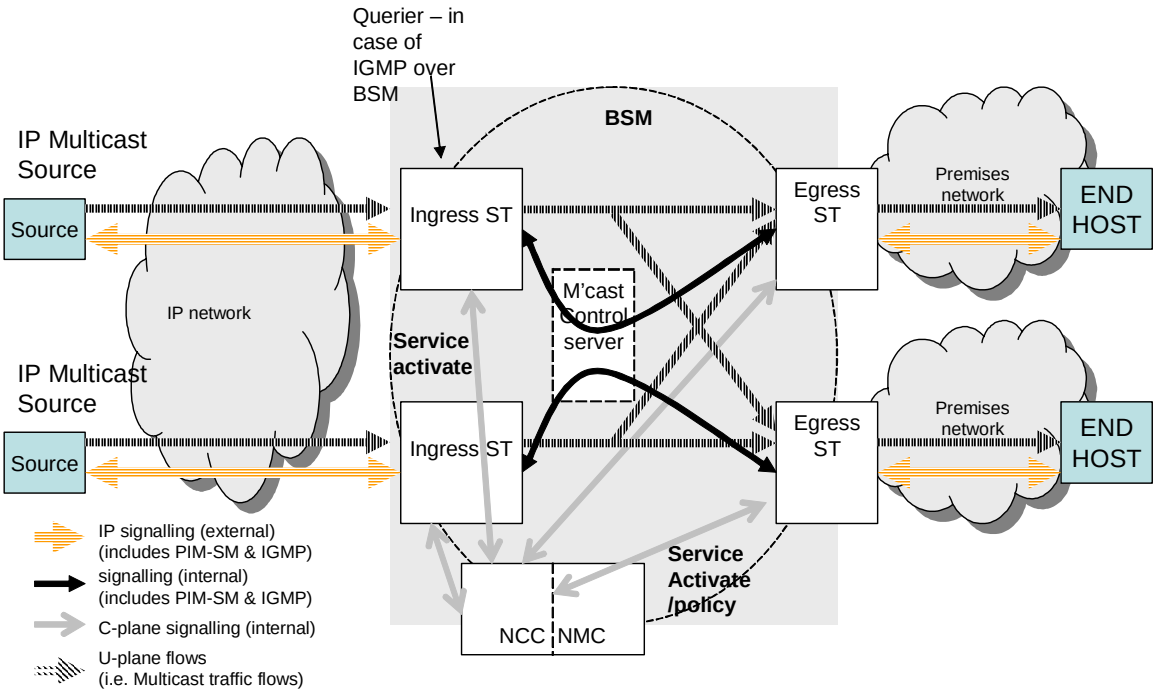

Figure 4.4: Mesh Pull Scenario

4.5.1 BSM Network Multicast Forwarding

Similarly to the Mesh Push case, there may be multiple ingress STs, which forward multicast flows from different sources and similar methods of separating the flows may be used as described in clause 4.4.1.

5 Control Plane Functional Architecture

This clause defines firstly the BSM Multicast Source Management Control Plane Functional Architecture and then defines the functional entities involved.

The BSM Multicast architecture shall be consistent with the definition given in [1].

The IGMP and PIM-SM protocols shall be implemented on BSM external interfaces as defined in [3], [4]and [5] respectively.

The Functional Architecture is based on the BSM network scenarios defined in clause 4. This architecture is focussed on the functional entities involved in the end-to-end BSM multicast control mechanisms that enable multicast flows to be forwarded or removed across the BSM from Ingress to Egress. The architecture must support dynamic control of multicast groups, allowing groups to be added and removed on demand.

BSM Multicast Source Management refers to the combination of Control Plane functions needed to create, maintain and remove BSM multicast distribution trees, and which includes Multicast Control Management (using PIM and IGMP), Multicast Access Control, and Multicast Address Management.

5.1 Architectures

The architecture is defined with three variants for the scenarios of clause 4 owing to the different functions and interactions between functions required, as follows:

1) Push (Star and Mesh).

2) Star Pull.

3) Mesh Pull.

These architectures are described below.

5.1.1 Push (Star and Mesh) architectures

The Push (Star and Mesh) architectures are shown diagrammatically in Figure 5.1.

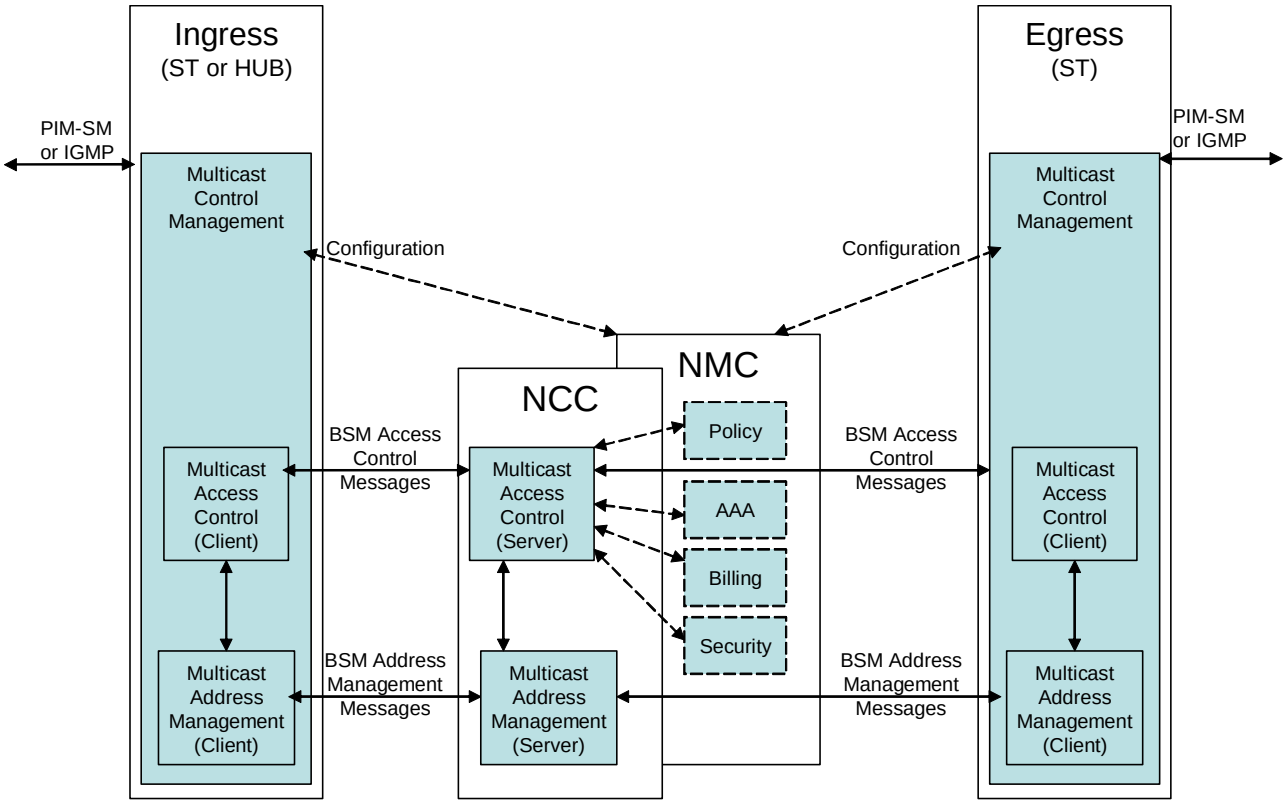

Figure 5.1: BSM Multicast Source Management Control Plane Architecture (Star and Mesh Push)

Figure 5.1 illustrates the relationship between the Multicast Control Management, Multicast Access Control, and Multicast Address Management functional entities and their elements located within the ingress/egress STs, the NCC/NMC. It also indicates the general paths of the messaging.

For all the scenarios the NCC is understood to be concerned with BSM SI-SAP and SD layer functions. The NMC is considered closely related to, or part of, the NCC, whose actions are performed under the aegis of the NMC for aspects such as policy, security and authentication. The functions of the NMC will not be detailed further in the present document, but, as usual, will be assumed to be present in the background.

In the Push scenarios of clause 4, no interaction at IP layer is allowed between ingress and egress, nor on the "vertical" interactions between IP layer and below. Instead the Access Control and Address Management are configured manually as shown in the figure.

5.1.2 Star Pull Architecture

The Star Pull version of the architecture is shown in Figure 5.2.

It differs from the Push architectures in the use of a layer 3 protocol between Ingress ST and Egress STs to control group membership dynamically, instead of by configuration. Also the Multicast Control Management Entity in the Ingress ST handles the interface to the Multicast Access Control.

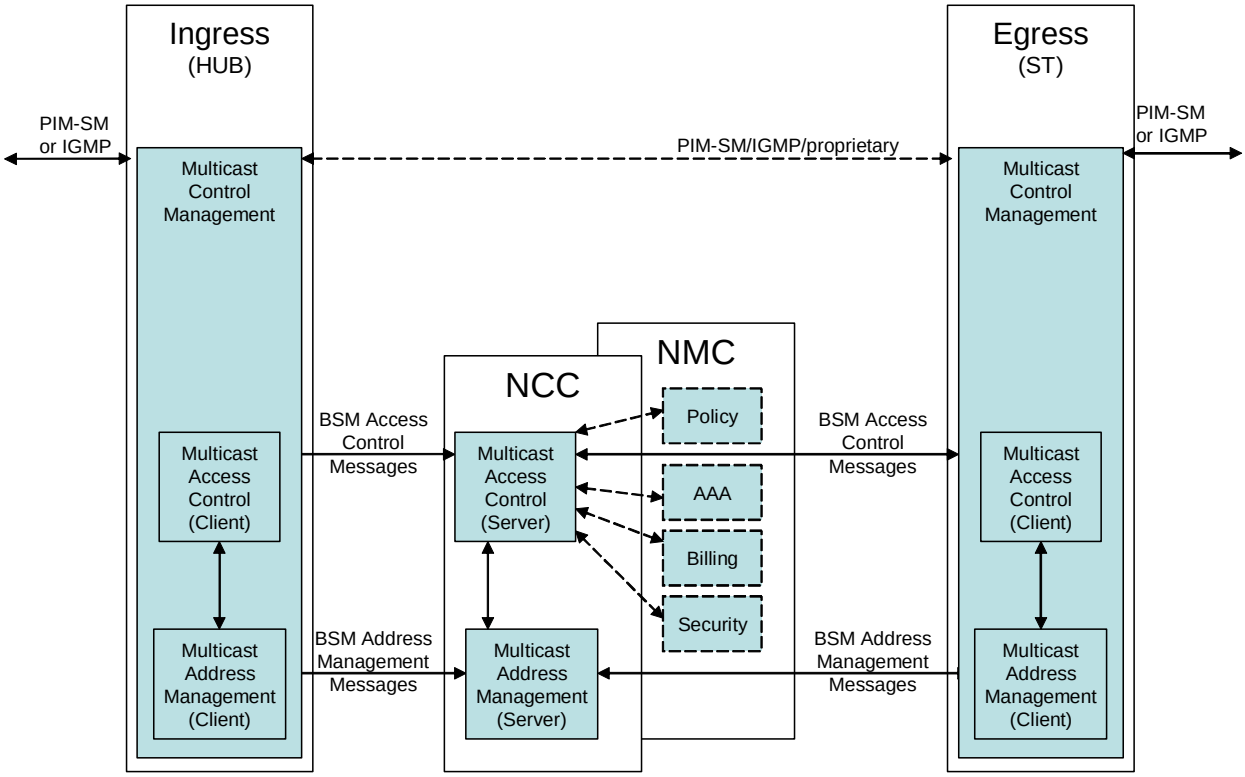

Figure 5.2: BSM Multicast Source Management Control Plane Architecture (Star Pull)

5.1.3 Mesh Pull Architecture

The Mesh Pull version of the architecture is shown in Figure 5.3.

It differs from the above architectures in the (optional) use of a BSM Multicast Control Management Server at layer 3 in order to make a star configuration of L3 control message paths (see clause 4.5).

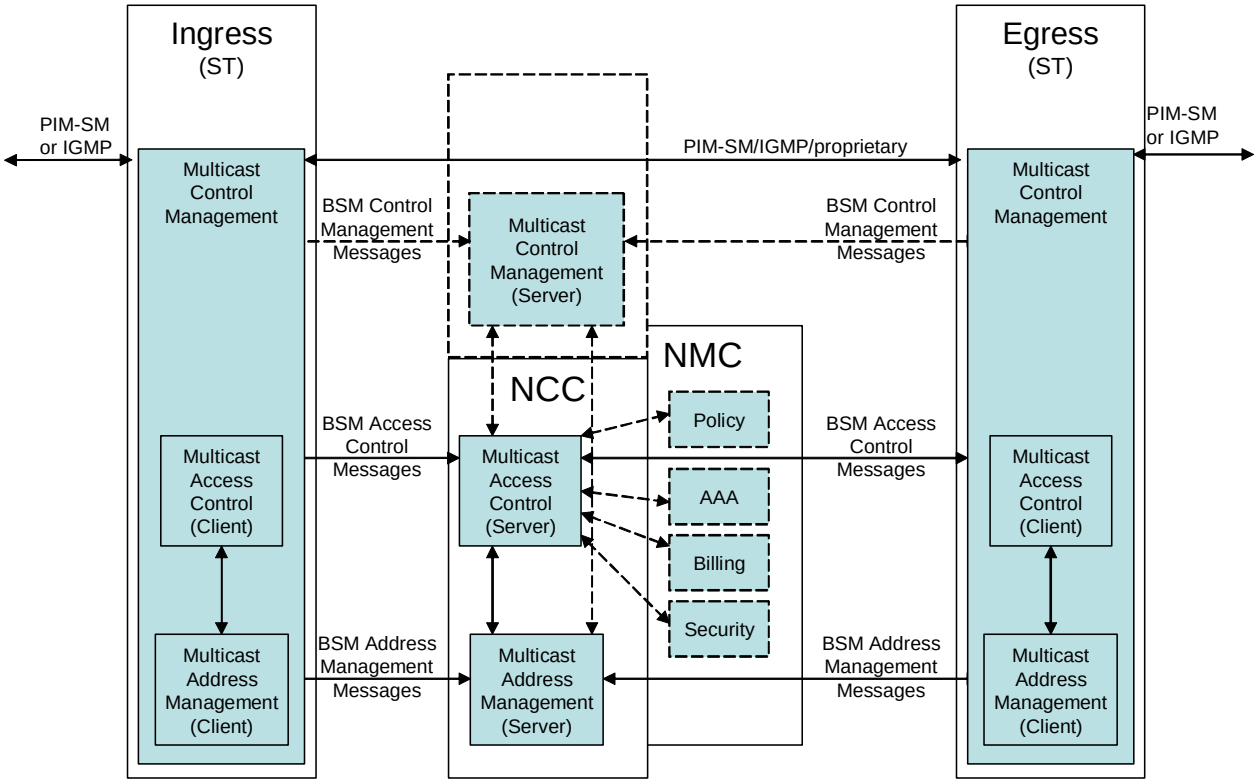

Figure 5.3: BSM Multicast Source Management Control Plane Architecture (Mesh Pull case)

The functional entities are defined in the following clauses taking the Mesh Pull case as the functional baseline.

5.2 Functional Definitions

5.2.1 BSM Multicast Control Management

BSM Multicast Control Management (MCM) refers to interworking and possibly adaptation functions for end-to-end IP PIM and IGMP multicast protocols; these functions may also involve snooping or proxying [L] and [M].

The MCM entities in the STs are also defined as the master functional entities responsible for invoking their sub-entities Multicast Access Control, and Multicast Address Management, and for overall coordination of these functions. The MCM in each ST, either ingress or egress, is responsible for the forwarding downstream from that ST and for all the implied procedures necessary.

In the Push and Star Pull scenarios, the MCM entities are located only in the STs and Hub.

In the Mesh Pull scenario, the MCM entities are located in the STs as clients (MCMC's) and are connected to a centralized MCM server (MCMS), which allows coordination of the IP-layer protocols across the BSM.

As indicated in clause 4.3 the BSM internal multicast control messages can be IGMP or PIM or even a non-IETF protocol (e.g. S-IGMP [E]).

NOTE: For the BSM it is often important to reduce the signalling overhead as far as possible and hence less garrulous adaptations of IETF protocols are desirable (e.g. S-IGMP).

Whichever of these protocols is chosen, it is mandatory that the ingress and egress ST's interwork with IP protocols externally. It is often cost-effective, however, to also employ IETF protocols internally wherever possible since compatible equipment can be more readily obtained (and hence S-IGMP retains IETF protocol compatibility at the client side).

5.2.1.1 Egress ST

IGMP or PIM is assumed as the protocol used to control group membership in the Egress ST's premises network. The ST can act in several ways:

1) As a local router with PIM on the BSM side and IGMP on the CPN side. The ST is an IGMP Querier.

2) As a PIM router with PIM on both sides (if there is at least another router in the CPN).

3) As an IGMP proxy (BSM-side client and CPN-side querier) with IGMP on both sides.

4) As an IGMP snooper with IGMP passing pseudo-transparently.

Group membership of the ST is determined by requests from attached hosts and from any downstream CPN devices.

5.2.1.2 Ingress ST

The Ingress STs can act in several ways:

1) As a PIM router with PIM on both sides (the STs are PIM routers).

2) As a local router with PIM on the Internet side and IGMP on the BSM side. The Hub act as an IGMP Querier.

3) As an IGMP proxy or snooper with IGMP on both sides.

For each new group membership request it receives on a particular BSM interface, the ingress ST's MCMC requests permission and resources from the BSM Multicast Access Control to construct, modify or remove a BSM distribution tree, for that group and for the relevant interfaces . If the response is positive, the MCMC subsequently adds the group to its list for the specific interface. and send a listen request to the SI-SAP for the instruct the allow forwarding of the associated multicast flows.

For any group membership expiry, the MCMC requests removal of the relevant resources from the BSM Multicast Access Control.

The Ingress ST must interact with the BSM Multicast Address Manager to obtain a BSM GID (in the case of IGMP over BSM) for forwarding.

5.2.1.3 PIM/IGMP Configurations

As indicated in clause 5.2 there are several options for the configuration of PIM and IGMP over the BSM. The following figures are based on those from TS 102 294 [1].

5.2.1.3.1 IGMP-over-BSM (PULL scenarios)

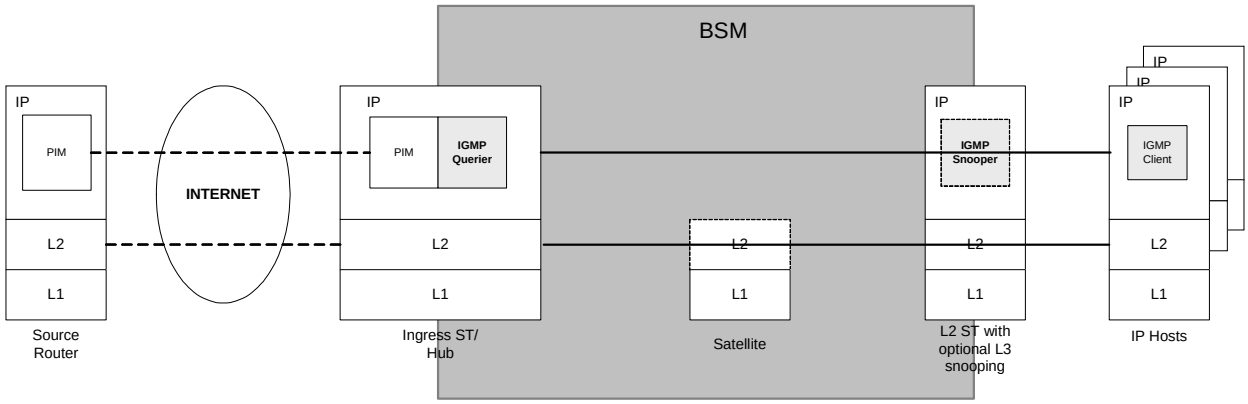

Figure 5.4: IGMP-over-BSM Architecture with Egress ST Layer 2 (Switching) and optional IGMP snooping

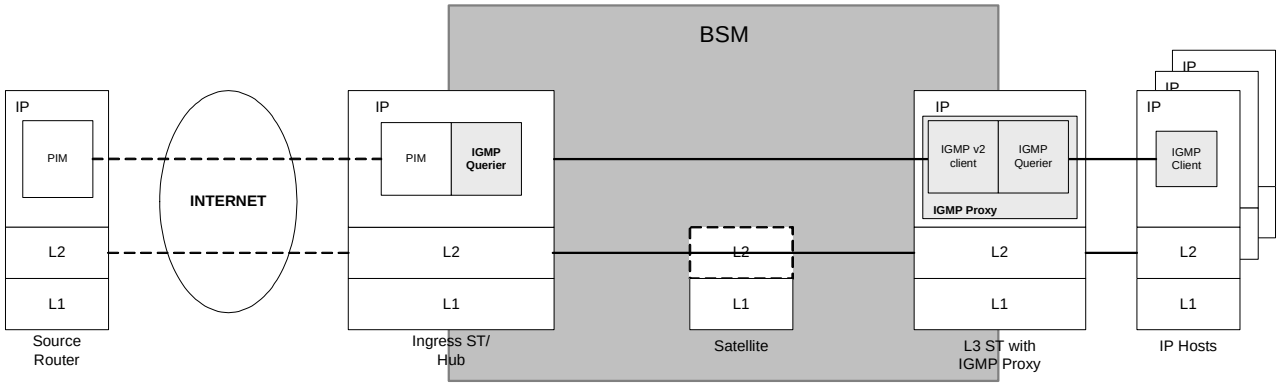

Figure 5.5: IGMP-over-BSM Architecture with Egress ST Layer 3 Forwarding

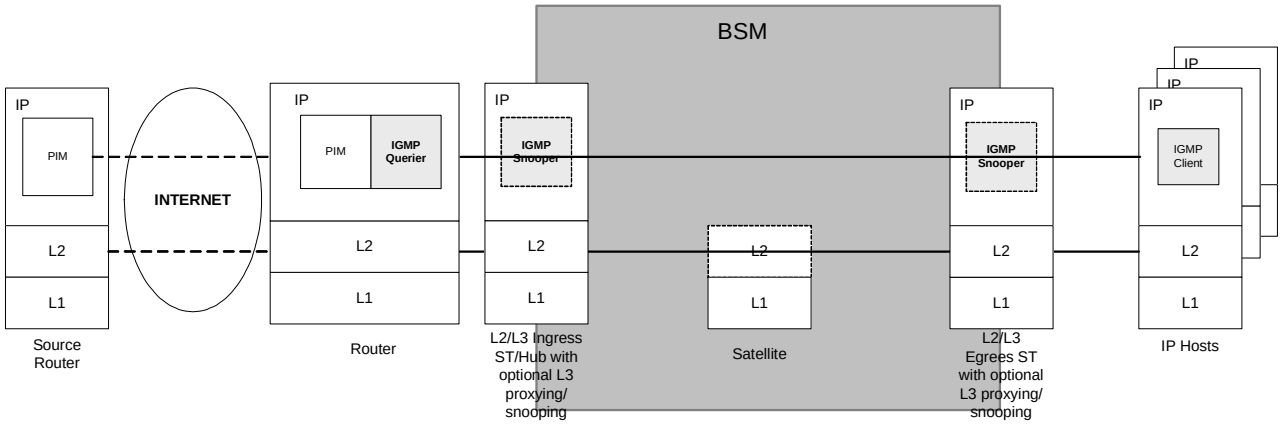

Figure 5.6: IGMP-over-BSM Architecture with IGMP external to BSM and optional IGMP proxying or snooping at Ingress and Egress

5.2.1.3.2 PIM-over-BSM (PULL scenarios)

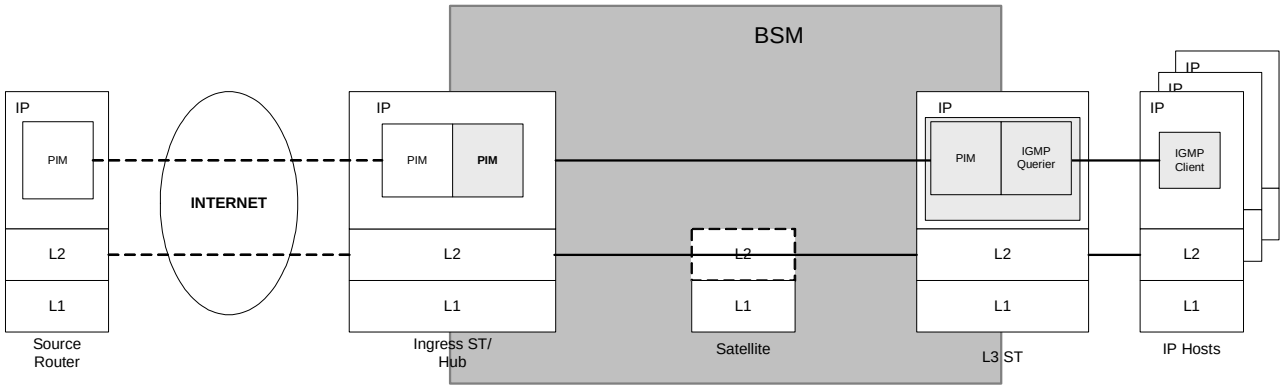

Figure 5.7: PIM-over-BSM Architecture with PIM external to Ingress and IGMP external to Egress STs

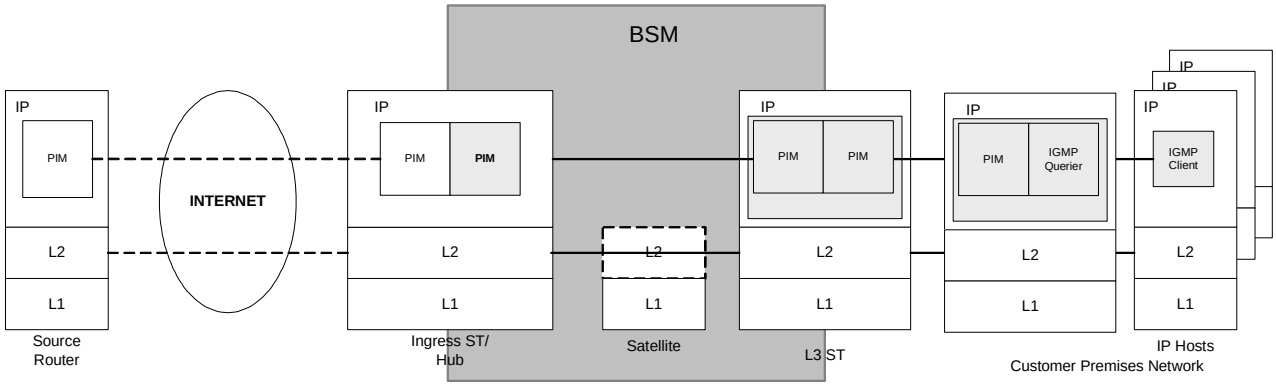

Figure 5.8: PIM-over-BSM Architecture with PIM external to STs

5.2.1.4 Multicast Control Management Server

The MCM server is of particular interest for the Mesh Pull case since for IGMP/PIM signalling, satellite network return links do not behave in the same way as LANs for which IGMP and PIM protocols were designed. Several cases can be made:

a) If there are several PIM-SM ingress routers on a LAN, one of them must be chosen to operate as the designated router (DR). The DR is then responsible for sending triggered Join/Prune and Register messages toward the RP. The DR election mechanism involves neighbouring routers sending PIM Hello messages to each other. The sender with the highest network layer address assumes the role of DR. Each router connected to the LAN sends the Hellos periodically.

   Downstream routers (e.g. BSM Egress STs) need to determine a single forwarding ST when the same (S,G) is available at multiple Ingress STs. Normally this if determined by the PIM Reverse Path Forwarding, (RPF) check to the Source (S), but there are circumstances where this might not be the case, e.g., when using multiple unicast routing protocols on that LAN, transient routing topology changes. Where the RPF check does not result in a single Ingress ST, the PIM Assert mechanism is used to determine the elected forwarder, and therefore where to send subsequent Joins.

   In the case of a BSM mesh network, there is not normally a point-multipoint channel provided to all such routers and no easy way to resolve these ingress conflicts without special measures. The DR election and Assert procedures may involve excessive mesh communications overhead between ingress STs as the Ingress routers are connected by satellite. The use of the MCM server could alleviate the election processes and reduce the signalling traffic, and allow the MCM server to act as a proxy for the DR.

b) In normal IGMP operation join/leave reports are sent using multicast addresses, but, as described in clause 4.5, it is usual for the satellite to offer only unicast channels rather than a multicast return channel to all relevant attached STs. The MCMS could then act as a proxy for the multicast destinations intended, and forward reports to the relevant ingress ST or other STs concerned.

This architecture requires configuration of control plane routing and/or modification of the PIM or IGMP protocols. The details are for further study.

There is a possibility to resolve this multiplication problem by use of PIM and "assert" messages between Ingress STs, but this would require additional mesh communications between Ingress STs, which should be avoided.

5.2.2 BSM Multicast Access Control

BSM Multicast Access Control (BMAC) refers to the functions that control the authorization and resources for BSM multicast distribution trees; i.e. creating and controlling the lower layer resources that are required for the forwarding of multicast data (e.g. for a given set of multicast groups in IGMP) through the BSM network from ingress ST/Hub to egress STs.

The BSM lower layer resources include the uplink resources to transmit the multicast data from the ingress ST/Hub, the satellite payload resources to relay and optionally replicate the multicast traffic (including possible satellite payload reconfiguration) to several beams or physical layer channels and the downlink resources (including optionally multiple downlink spot beams) to transmit the multicast traffic to the egress STs.

NOTE 1: BMAC is proposed as a new function that is not defined in TS 102 294 [1] (V1.1.1). It should be added in the next revision of that document.

Access control of multicast services, as for unicast access control, can be operated on a guaranteed service or relative service basis, and which can be related to IP multicast using either the IntServ per-flow model or using aggregated flows (e.g. DiffServ) respectively.

The guaranteed service solution, however, has potential scalability problems for many flows where an SD bearer has to be handled for each one.

The relative service solution avoids scalability problems and can be used to manage an aggregate of multicast flows (in a similar way to unicast) [J]. In this case the BMAC allocates BSM resources (QIDs) to multicast flow aggregates, and may offer different classes of service. A Service Level Agreement (SLA) is established between the BSM and the external network operator on the resources and service classes expected. Multicast flows are classified at the Ingress to the BSM and are allocated to the appropriate IP queue and subsequently to one or more QIDs. Excess traffic is treated according to the SLA. In this way IP group membership can be dynamic and scalable, and can also conform to resources available. The aggregate flows can be treated simply on a best effort basis, or can be prioritized into classes of service (e.g. as in the DiffServ model).

NOTE 2: There is a potential problem with multicast over DiffServ, see [J]. Hence a more general approach to classes or priorities of aggregated multicast services may be more appropriate.

5.2.2.1 Resource Management

The BMAC should be able to interact with similar resource management functions as for unicast, as defined for IntServ and DiffServ in [G] and [H] respectively. The generalized relationship between the BMAC, MCM and resource management functions in an Ingress ST is shown in Figure 5.9. This diagram indicates options for IntServ and DiffServ. The resource management can be static (configured) or dynamic (via signalling).

The unicast functional architectures for IntServ and DiffServ resource management are given in [G] and [H] respectively. The differences for multicast are that the Control Plane has inputs from the MCM and BMAC, and the user plane forwards flows to relatively few BSM GIDs rather than BSM UIDs. These architectures for BSM QoS processing apply primarily to the Ingress STs, but a similar scheme can be applied to the Egress ST where resources in this case are mainly under the control of the attached network.

5.2.2.2 Egress ST

The egress ST's Multicast Access Control Client (MACC) receives messages from the Multicast Access Control Server (MACS) to open or close SI-SAP resources for GIDs, and messages concerning policy regarding onwards forwarding.

5.2.2.3 Ingress ST

The ingress ST MACC receives requests for or cancellations of resources from the MCM, and interacts with the MACS to confirm or deny the request.

The Ingress ST must interact with the BSM Multicast Address Manager to obtain a BSM GID that is used when forwarding the multicast data via the U-plane.

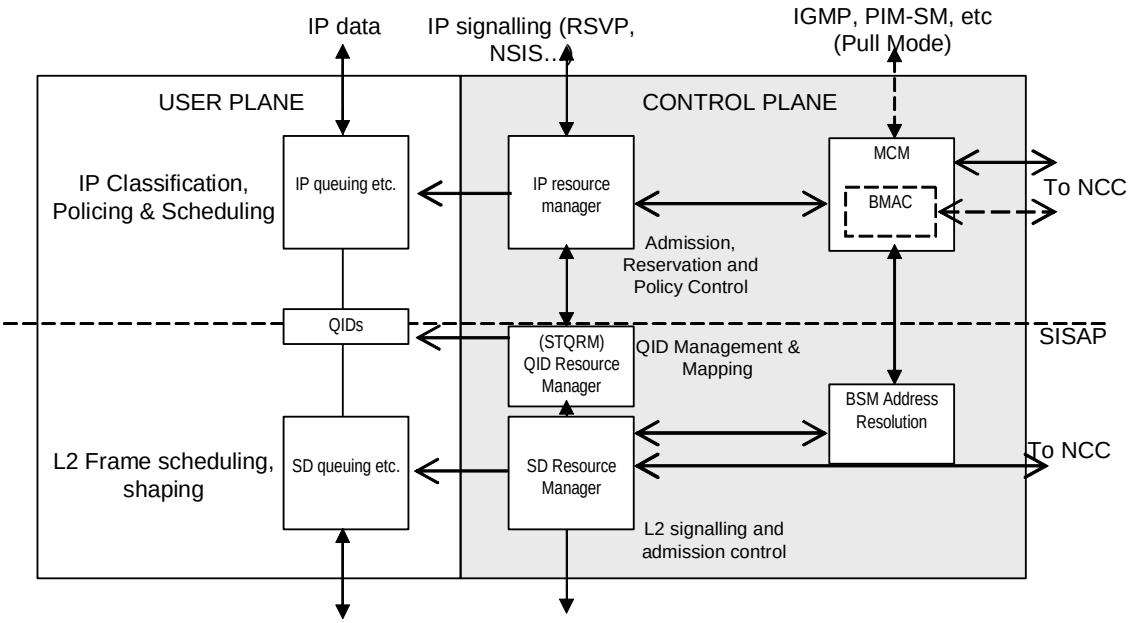

Figure 5.9: Ingress ST multicast resource management functional relationships

5.2.2.4 NCC

The MACS receives messages from MACCs (or from central management in Push mode) to open or close SI-SAP resources for GIDs.

If it decides to grant resources, it also requests the Multicast Address Management server (MAMS) to allocate a GID to the IP group.

It issues commands to the ingress and egress STs (in a given beam/channel), and satellite if necessary, as a result.

It informs relevant egress STs on policy on which multicast groups may be forwarded to the attached networks.

5.2.3 BSM Multicast Address Management

Multicast Address Management refers to the general functions of multicast address resolution between multicast IP groups and BSM GIDs, as well as between a BSM_GID and a lower layer identity. The function also includes interacting with the Multicast Access Control and the Multicast Control Manager.

The NCC or the Ingress ST initiates requests to the MAMS for address resolution, if the address association is not already known.

The egress ST's MAMC (Multicast Address Management client) may send requests to the MAMS for GID-to-IP group association.

The MAMS either issues explicit allocations to ingress and egress MAMCs, or sends tables periodically.

Once the association is known, and forwarding is enabled by the BMAC, the Ingress ST forwards flows received on the specified group address to the SI-SAP interface(s).

5.2.3.1 BSM Multicast Address Resolution

Multicast Address Resolution (MAR) includes functions both above and below the SI-SAP. Above the SI-SAP it is necessary to associate a multicast IP address [P] to a BSM_GID. Below the SI-SAP it is necessary to associate a BSM_GID to a lower layer address, typically a MAC address. This function is similar to the unicast Address Resolution function defined in TS 102 460 [F] but has some important differences:

a) The BSM_GID is an SI-SAP group identifier. Each BSM_GID may have associated with it a group of STs that want to receive this multicast flow.

b) SI-SAP Address Scoping may be needed to separate multicast flows with overlapping IP multicast addresses. Some recommended scoping techniques are defined in annex A.

c) Multiple IP multicast addresses may be associated with the same GID. For example, if the default address mapping based on RFC 1112 [3] is used, then 32 IP multicast addresses can be mapped to each GID.

d) AR tables are assumed to be preconfigured by the Address Management system before the multicast data is allowed to flow (in or out) through the BSM System. These tables can be either static tables (configured or rule based tables) or dynamically populated tables using the B-AR function as defined in TS 102 460 [F].

The BSM Architecture must provide a service where MAR is supported for MAR clients in all STs by a central multicast address resolution server. The MAR server could, in principle, be located anywhere but it is realistic to assume that is under the control of the BSM operator since it needs knowledge of the BSM address space. Typically, the MAR Server will be located at a Gateway or at the NCC. Having the MAR server located at the NCC may be appropriate if the MAR function is used to support Access Control, i.e. allowing and denying multicast access to the BSM network.

MAR is a C-plane function. However, two distinct processes are required for MAR to function. Above the SI-SAP, a BSM_GID must be associated with a multicast IP address. Below the SI-SAP, a BSM_GID is associated with a MAC address or equivalent lower layer identity. MAR may be used whenever a multicast packet of a new group is to be forwarded across a BSM network [Q].

6 Source Management Protocol Architecture

This clause illustrates the specific relationships of Multicast Source Management functional entities outlined in clause 5, and expands on the multicast protocol architecture and in TS 102 294 [1]. It considers the three constituent functions of Multicast Source Management:

1) Control Management (MCM).

2) Access Control (BMAC).

3) Address Management (MAM).

The protocol architecture depends on whether the ST is at the Ingress or the Egress.

At the Ingress the ST negotiates multicast resources and forwarding over the BSM, but the reception of multicast groups is on the attached network.

At the Egress the ST does not need to negotiate BSM resources but needs to filter BSM GIDs for reception of groups.

The protocol diagrams and functional definitions below assume the full functional cases of Pull scenarios unless otherwise stated. Push cases can be derived by elimination of dynamic group membership and associated functions.

The multicast User Plane is illustrated in each case in one direction - from Ingress to Egress. The Control plane is shown with bidirectional message flows at IP layer. The unicast User Plane is not shown for clarity.

6.1 Ingress ST

Figure 6.1 shows the multicast protocol architecture of the Ingress ST.

The IGMP/PIM protocol entity (in the MCMC) establishes the IP group membership list (under the aegis of the BMAC) for each of the Ingress ST BSM interfaces. Whenever there is a change of aggregate group membership over all of these interfaces, and/or periodically as necessary, the MCMC sends a resolution request for any new groups to the lower layers of the attached network in order to obtain associated link layer addresses. It also sends a resolution request for the groups to the SI-SAP to obtain associated GIDs on the BSM side. Reception and forwarding of multicast groups is controlled by the MCMC, and having obtained the BSM resources necessary, the MCMC issues a "Listen" command to the attached network interface together with the binding of relevant IP groups and multicast link layer addresses. The MCMC also issues a "forward" command to the IP forwarding engine together with the binding of the groups to GIDs.

NOTE: If The Ingress ST has several BSM interfaces with separate group memberships, each interface must be identified with a different IP address.

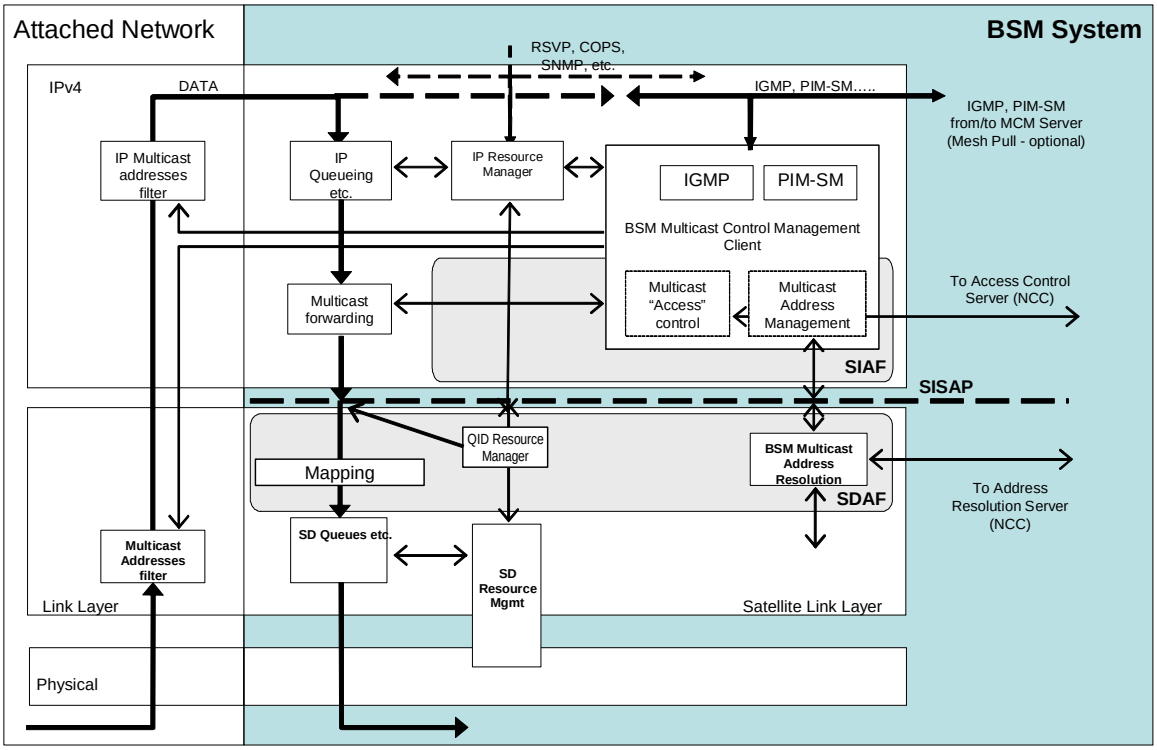

Figure 6.1: BSM Multicast Ingress ST protocol stack

6.2 Egress ST

Figure 6.2 shows the multicast protocol architecture of the Egress ST.

The IGMP/PIM protocol entity (in the MCMC) establishes the IP group membership list (under the aegis of the BMAC). Whenever there is a change of group membership it issues a join request to the upstream router.

The MCMC also sends a resolution request for any new groups to the SI-SAP to obtain associated GIDs on the BSM side. It also sends a resolution request for the groups to the lower layers of the attached network in order to obtain associated link layer addresses. Reception and forwarding of multicast groups is controlled by the MCMC, and it issues a "Listen" command to the IP forwarding engine together with the binding of the groups to GIDs. The MCMC also issues a "forward" command to the attached network interface together with the binding of relevant IP groups and multicast link layer addresses.

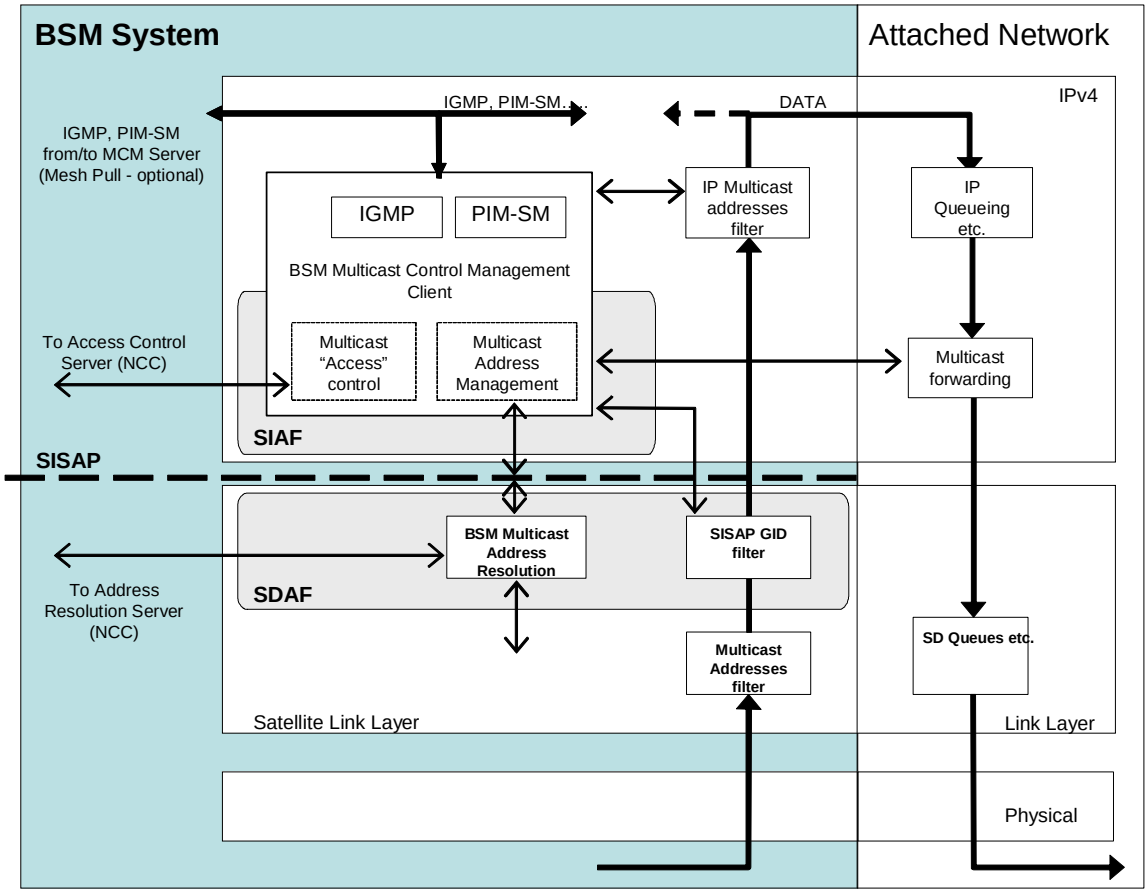

Figure 6.2: BSM Multicast Egress ST protocol stack

6.3 Forwarding from the Ingress ST

6.3.1 User plane - Forward Links

As explained in clause 4.1, a satellite system may split its overall coverage into several parts by means of multiple beams, by physical layer channels or by link layer channels. Whichever of these is used, they can be treated as different paths over which multicast data can be transmitted.

At the IP layer the different paths are treated as different output interfaces to the Ingress ST, similar to a standard multicast router. The Ingress ST would maintain group membership for each of its output interfaces. The output interfaces are included in the ST's forwarding table and group flows are forwarded to the appropriate interfaces. This is illustrated in Figure 6.3.

Each output interface is identified by a combination of QID and GID for a given instance of an SI-SAP. If necessary separate interfaces can be further differentiated by multiple instances of the SI-SAP.

Forwarding to each instance is controlled by its group membership.

Each instance needs to be associated with a physical layer and/or link layer channel.

6.3.2 Control Plane - Return Links

For the purposes of group membership, there is a need to associate return signalling channels from egress STs with multicast forwarding output interfaces as defined above.

There are different options for return links as described in clause 4.5 and as illustrated in Figure 6.3. The BSM SI-SAP needs to take into account these options.

The details of this scheme are out of scope of the present document.

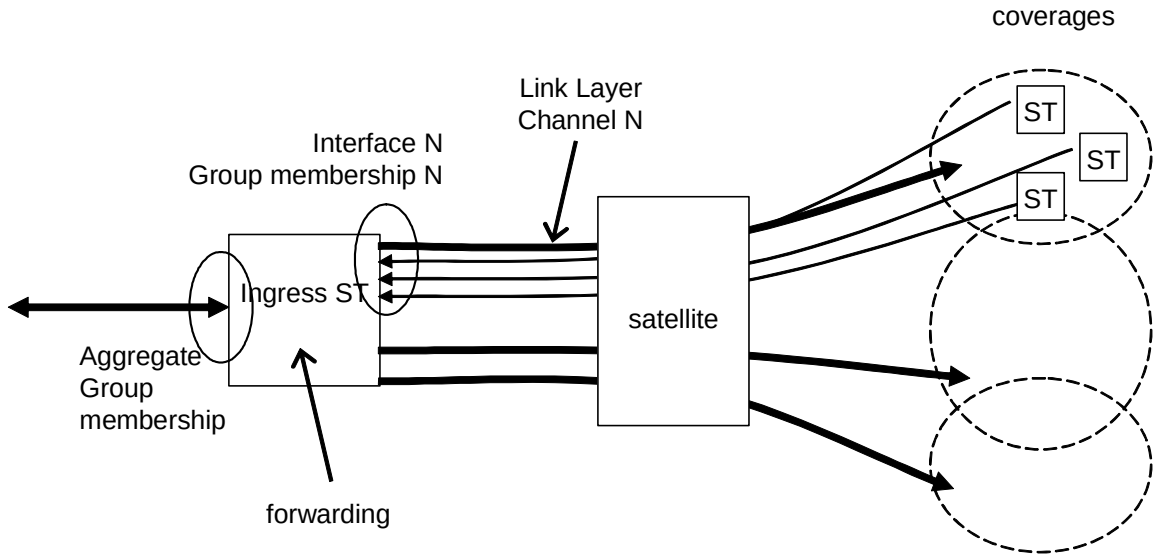

Figure 6.3: Forwarding and PIM join/prune or IGMP reporting scheme

7 BSM Multicast Protocol Message Sequence Charts

The following diagrams are shown for a given SI-SAP interface instance from the Ingress ST towards the BSM. As described in clause 6.3, there may be several such interfaces and a separate protocol operation is required for each one.

7.1 Push

Figure 7.1 is valid for both Star and Mesh Push cases.

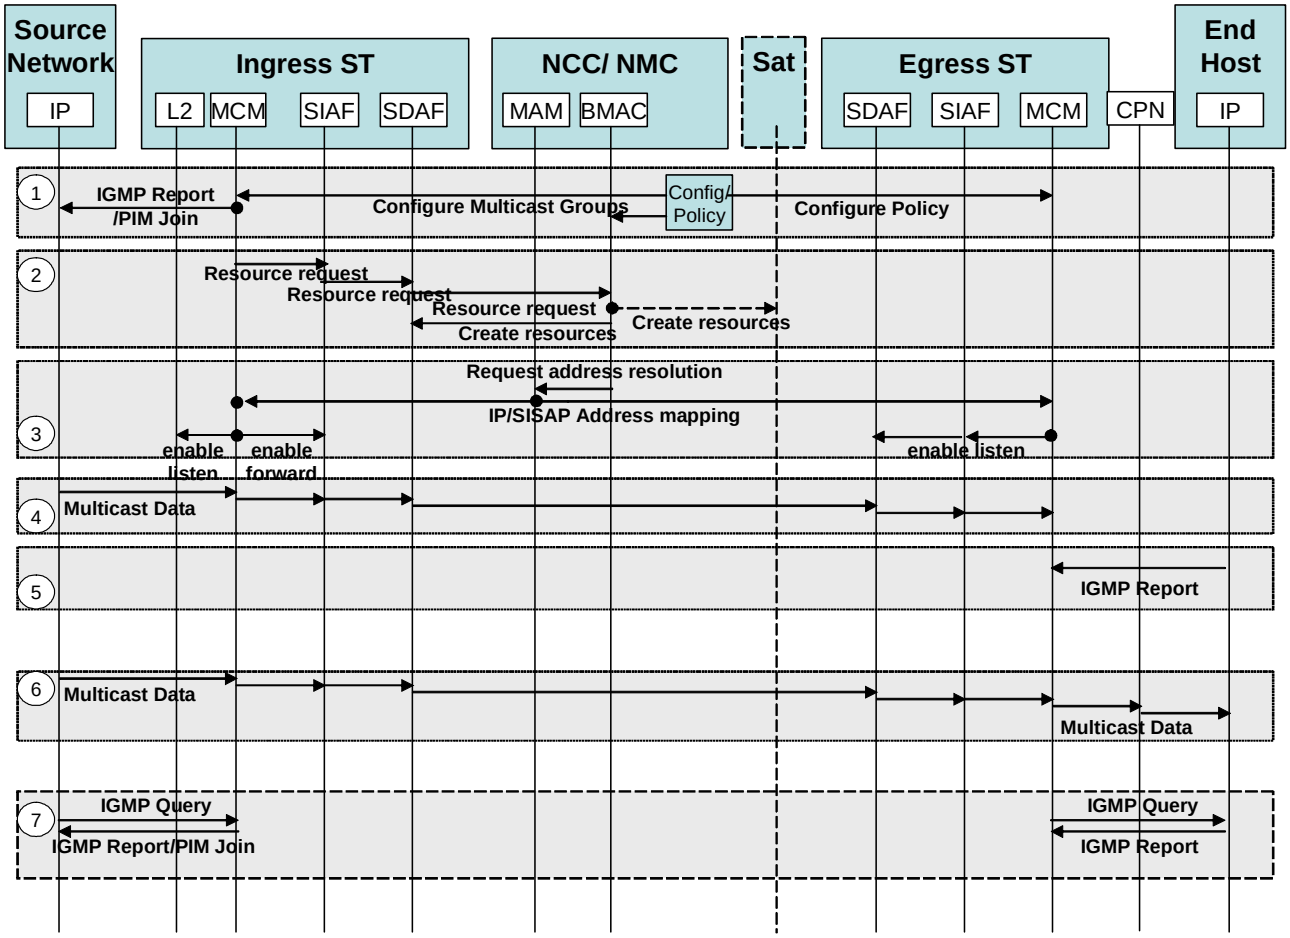

Figure 7.1: Push Scenario Messages

In this scenario the following steps are shown:

STEP	DESCRIPTION
1	The NCC/NMC configures the BSM to become a member of an IP multicast group, by issuing a group join command to the Ingress ST. The Ingress ST starts normal PIM/IGMP operation by sending a join request etc. Note that any group that is joined at this stage is not used by the Ingress ST until step 3. The NCC/NMC also authorizes the BMAC server to create BSM resources, and configures the forwarding policy in the egress STs.
2	If the Ingress ST MCM adds a group to its membership list, it may if necessary request resources for this group (e.g. via the QoS resource manager) by issuing a request to the BMAC server. If the satellite is an OBP type, a resource request may also be issued to it.
3	The BMAC server coordinates Address Resolution within the BSM by issuing, if necessary, a request to the MAR server which associates a GID with the IP group address. The BMAC server then issues the results to the MCM of the Ingress and Egress STs. The Ingress MCM then issues a listen command to the network interface, and a forwarding command to the ST BSM-side engine. The Egress MCM's issue listen commands to the SI-SAP interface.
4	Multicast data with the specified group address arriving at the ingress is forwarded over the BSM up to the Egress STs,
5	Unsolicited IGMP "join" reports or PIM joins may be sent by hosts to the Egress ST at any time, or they may be sent in response to a "Query".
6	The Egress ST begins forwarding multicast data for the group.
7	Further IGMP Queries and reports may be sent as necessary at any time at both ingress and egress sides.

As in all these scenarios, the above sequence is repeated if another IP multicast group is added, and it also applies in a similar fashion if an IP multicast group is removed from the BSM system.

7.2 Star Pull

Figure 7.2 shows the message flows for the Star Pull scenario.

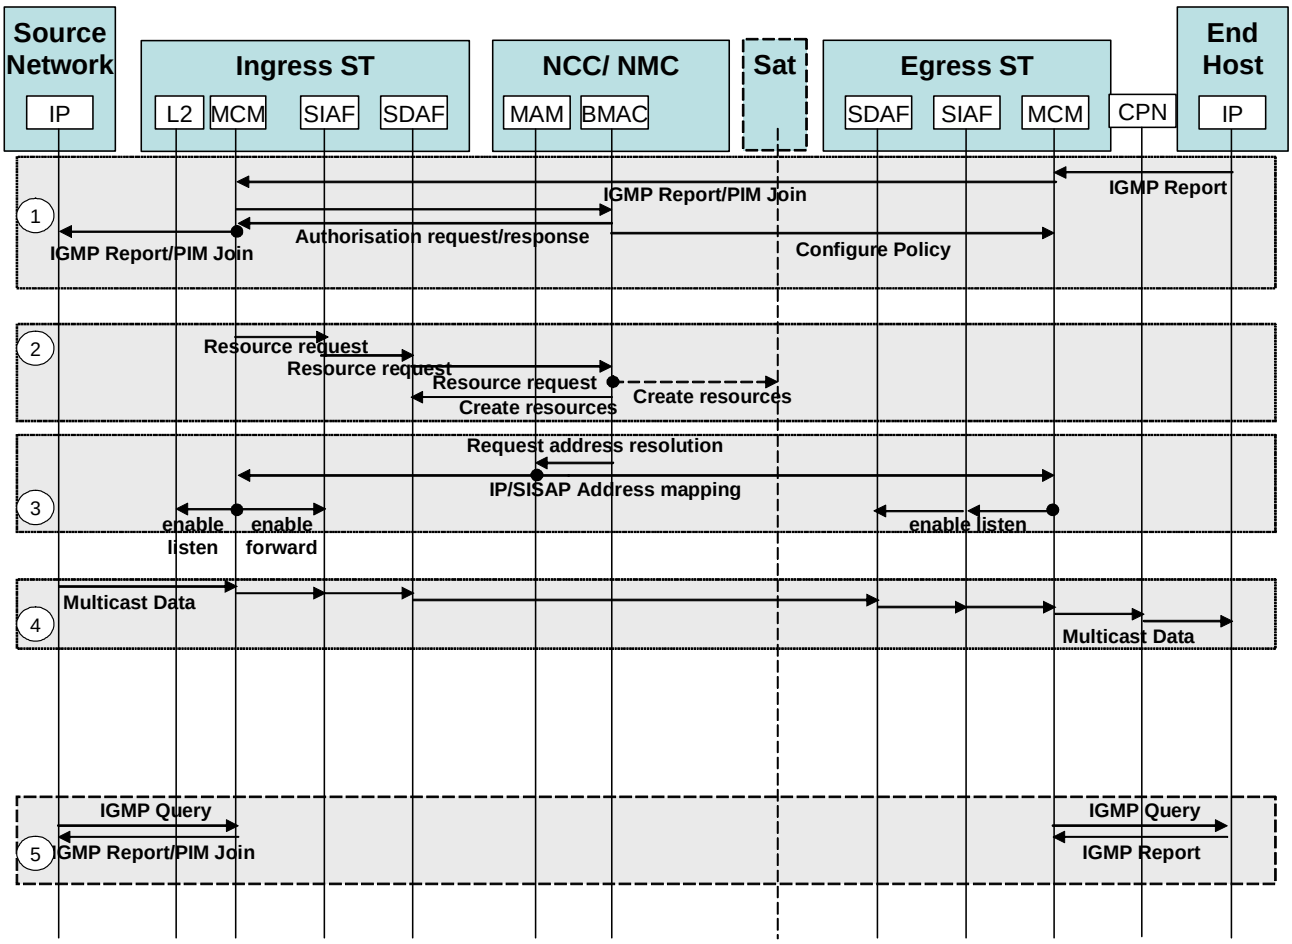

Figure 7.2: Star Pull Scenario Messages

STEP	DESCRIPTION
1	In response to an unsolicited IGMP "join" report or PIM join sent by a host to the Egress ST, or in response to a "Query", the egress ST (MCM) adds a new group to its membership list. It then issues a similar join message upstream to the Ingress ST. If this is a new multicast group for the Ingress ST then the MCM requests authorization from the BMAC server. If authorization is granted, the MCM adds the group to its membership list and proceeds subsequently as for the Push case above. Note that the IGMP protocol over the BSM may be adapted as described in clause 5.2.1 to reduce the signalling overhead.
2 to 5	The subsequent steps are controlled by the NCC/NMC and are identical for the Push case above. If successful multicast forwarding can begin immediately after step 3.

7.3 Mesh Pull

Figure 7.3 shows the message flows for the Mesh Pull scenario. This detailed message sequence chart describes the baseline option of several options discussed in clause 5.2.1.4.

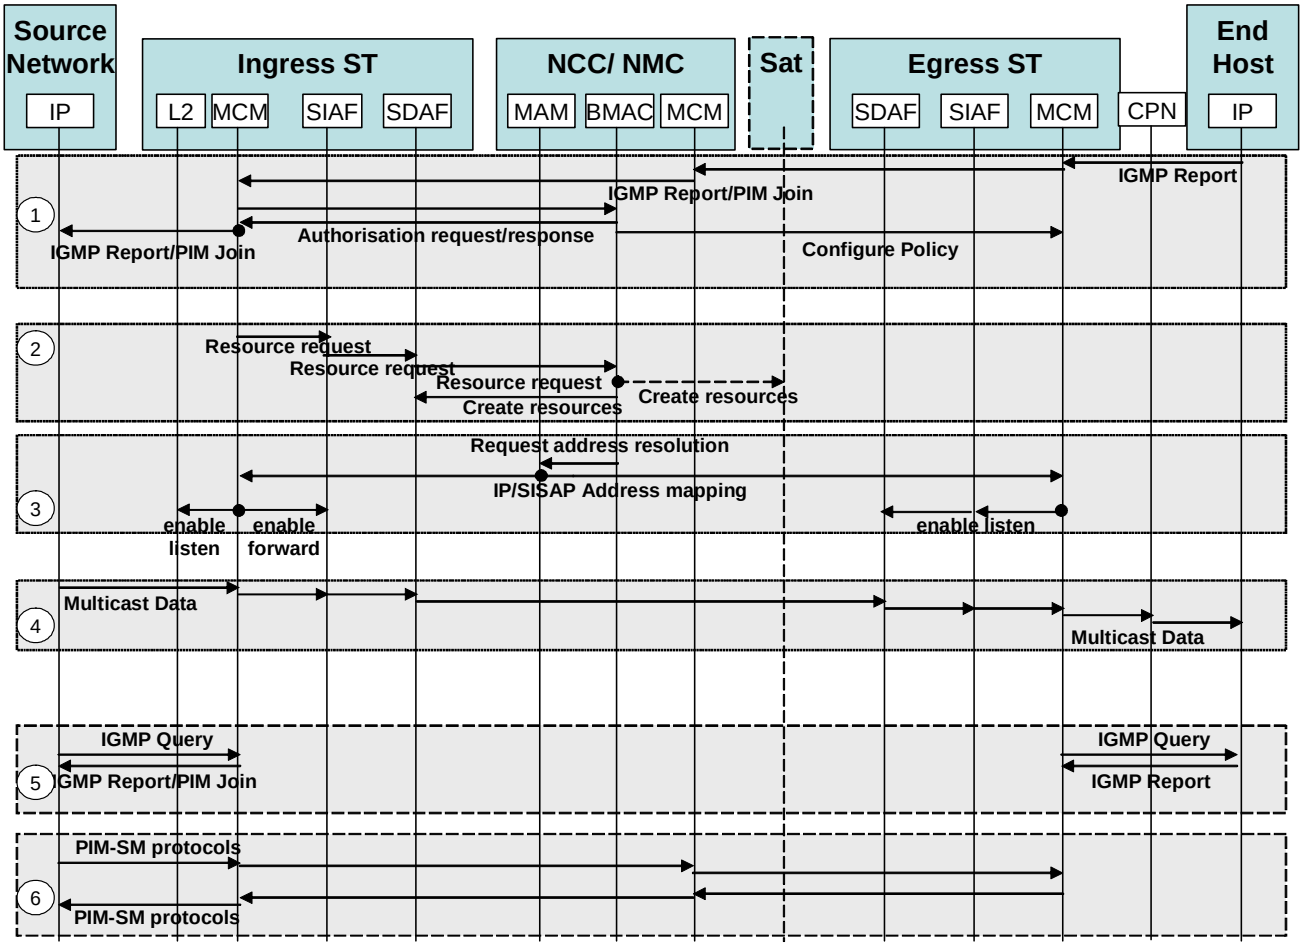

Figure 7.3: Mesh Pull Scenario Messages

The primary difference here from the Star Pull case is the (optional) introduction of the MCM server.

STEP	DESCRIPTION
1	In response to an unsolicited IGMP "join" report or PIM join sent by a host to the Egress ST, or in response to a "Query", the egress ST (MCM) adds a new group to its membership list. It then issues a similar join message upstream; in this example this message is sent to the central MCM server. The MCM server acts upon the message and decides to which Ingress ST to pass the message. If this is a new multicast group for the Ingress ST then the MCM requests authorization from the BMAC server. If authorization is granted, the MCM adds the group to its membership list and proceeds subsequently as for the Star Push case above.
2 to 5	These steps are controlled by the NCC/NMC and are identical for the Star Push case above. If successful multicast forwarding can begin immediately after step 3.
6	Any subsequent signalling messages between the Egress and Ingress STs are routed via the MCM server, including any other protocols (e.g. Hello, Assert) associated with PIM-SM networking.

Annex A (informative): Scoping of IP Multicast Addresses

A.1 Introduction

The BSM network architecture must be capable of supporting multiple "private" networks, including networks that use overlapping private IP address spaces. RFC 4259 [R] discusses this problem for MPEG networks: similar arguments apply to the general BSM case.

This annex first defines the IP multicast scoping requirements and then defines some recommended methods of supporting these requirements using the BSM SI-SAP architecture, in particular methods of address scoping using the SI-SAP and the BSM GID.

NOTE: This annex describes methods of SI-SAP address scoping to provide scoping (separation) of IP multicast traffic. These same methods of SI-SAP address scoping can also be used to provide scoping/separation of IP unicast traffic.

A.2 Requirements

A.2.1 IP multicast address scoping

IP multicast address scoping is the functionality that enables a BSM Network to support multiple "private" networks, each with a "private" scope whereby all the "private" multicast messages are localized to their own network domain and do not leak from one domain to another.

At the same time each "private" network should also have access to the global internet, including having access to "global" multicast messages; i.e. each network must simultaneously have the ability to receive multicast from globally routable addresses (remote sources).

In summary, the following requirements are assumed to apply for multicast scoping of "private" networks (see also draft-ietf-ipdvb-ar-xx.txt [Q]):

a) Private networks should be able to send and receive locally scoped (private) multicast (denoted by a local scope IP address, value or an addresses from the Scoped Multicast Address range RFC 2365 [W] and RFC 2375 [X]) within their private networks via a BSM network, including multicast between geographically separate stubs attached to different STs.

b) Each private area network may independently decide to receive globally routable multicast from globally routable addresses (remote sources).

c) Each private area may also receive globally routable multicast from sources within their own area; i.e. multicast with a locally scoped (private) source address that is not a globally routable unicast address.

NOTE: The combination of b) and c) means that a host may see both external global ASM {S,G} sources and internal local sources for the same G.

A.2.2 Multicast address resolution

Multicast address scoping requires each different "private" network to have a unique range of BSM GIDs (a different set of SI-SAP multicast group identities) within that BSM network. Given that the private domains may use the same IP addresses, this means, that the BSM architecture must permit alternative or flexible mappings between IP Multicast Addresses and BSM_GIDs.

The default mapping between IP Multicast Addresses and BSM_GIDs uses the standard IANA Ethernet address mapping defined in clause A.3.2. But to solve the multicast address scoping problem we also need to support the more general case with other mappings.

The BSM Multicast Address Resolution function has been defined to support a flexible mapping between Multicast IP Addresses and GIDs. BSM Multicast Address Resolution must be capable of supporting any valid mapping, including any mix of non-standard mappings in parallel with (or instead of) the default IANA-based mapping. Multicast Address Resolution can be implemented either using static procedures (e.g. pre-defined tables, containing the scoped mappings) or using dynamic procedures (i.e. using a multicast enabled SI-SAP address resolution function).

A.2.3 Network Address Translation (NAT)

With regard to Network Address Translation (NAT), there is a requirement to provide NAT for the unicast source address in multicast packets. This can be provided using a standard NAT function in the IP layer as defined in TS 102 460 [F].

The main scenario is the case where the source is in the private network (e.g. BSM) and is sending to the public network, then the "private" unicast source address needs to be translated.

There are also issues regarding the handling of the rendezvous point address that have to be considered, but this again is a unicast address translation.

The destination multicast address will normally be forwarded without translation since multicast address are receive-only and do not need to be (and should not be) private addresses.

A subset of BSM users may wish to support multicast between their "privately" scoped networks. This may require more complex address translation, but the problem is not specific to BSM and hence this requirement is considered out of scope for BSM.

A.3 OUI-based methods of SI-SAP address scoping

A.3.1 BSM_GID format

The BSM GID format uses the standard 48-bit universal IEEE 802 [Y] address format, which comprises two main fields:

- The first half (i.e. octets 0, 1 and 2) correspond to the Organizationally Unique Identifier (OUI) as assigned by IEEE, except that the assignee may set the LSB of the first octet to 1 for group addresses or set it to 0 for individual addresses. A second bit enables this universally administered (IEEE assigned) OUI to be replaced with a locally administered OUI.
- The second half, comprising the remaining 24 bits, (i.e. octets 3, 4 and 5) is administered by the assignee. This field can be used to map individual multicast addresses as defined below.

A.3.2 Default mapping for BSM_GIDs

The default mapping between IP Multicast Addresses and BSM_GIDs uses the standard IANA mapping for Ethernet multicast addresses as defined in RFC 1112 [3] as follows:

- a default multicast OUI using the IANA Ethernet multicast OUI (01-00-5E hex); and
- a default mapping of individual multicast group addresses, by placing the low-order 23-bits of the IP multicast address into the low-order 23 bits of the BSM_GID, with bit 24 set to '0'.

NOTE: As reported in RFC 1112 [3] there are 28 significant bits in an IP host group address, and hence this default mapping means that 5-bits of the IP address is lost and hence 32 different IP host group address are mapped to the same Ethernet multicast address.

This default case is the BSM equivalent of the standard Ethernet mapping. But in order to solve the multicast address scoping problem we also need to support other mappings for both fields; i.e. to support the use of both multiple OUIs and/or other mappings for individual multicast addresses.

A.3.3 Common OUI based scoping

A BSM Network may use a combination of locally administered (U/L='1') and universally administered (U/L='0') OUIs. Any universally administered OUI must be assigned with the IEEE registrar to avoid conflicts with IPX, AT, etc. As a result, a BSM network is only expected to have a small number of universally administered multicast OUIs, typically only one universally administered OUI per satellite network.

NOTE: The same OUIs that are used for multicast (for GIDs) may also be used for unicast (for UIDs). Equally, multicast and unicast may use different OUIs.

SI-SAP address scoping can still be achieved with a single common OUI by assigning a subset of BSM_GIDs to each network and applying filtering/ mapping in the SI-Layer.

However, more flexible SI-SAP address scoping can be provided by using multiple OUIs for the BSM_GIDs. A combination one (or more) common OUIs can be used in parallel with multiple private OUIs to provide one (or more) common (shared) multicast distribution network that operates in parallel with a set of private networks that each have a private OUI.

In this latter case, the BSM network could use a common OUI (e.g. a universally administered OUI) to provide a default IANA (RFC 1112 [3]) mapping for all globally addressed groups, thus permitting these groups to be sent to all STs (all multicast receivers) or even a selected subset of selected private networks. This would permit multicast messages with global scope (e.g. a CCN news feed) to be delivered to all private networks using a single transmission (i.e. the same internal BSM message). Without the use of this common OUI, the BSM network would have to replicate each multicast packet to every network that wants it.

A.3.4 Private OUI based scoping

Private OUI based SI-SAP Address Scoping is achieved by assigning a different OUI for each "private" network domain. Each network can then independently use the lower 24 bits (i.e. octets 3, 4 and 5) to assign GIDs and hence (for example) each network could continue to use the RFC 1112 [3] mapping for the lower 3 bytes of the GID.

The recommended method of using multiple private OUIs is use locally administered addresses; i.e. to set the Universally or Locally administered (U/L) address bit to '1'. Local administered addressing means that the entire address (i.e. 48 bits) has been locally administered. The BSM Network operator is free to assign any number of different OUIs (all with U/L='1') without requiring any co-ordination since the BSM Network can be operated as a closed network (i.e. the BSM_IDs and GIDs are only used by terminals attached to that BSM network).

A.3.5 Support for IPv6

The present document is limited to consideration of IPv4 addresses. However, a similar approach is expected to be used for IPv6 addresses.

An Ethernet mapping is specified for IPv6 multicast packets to an IANA-assigned OUI for IPv6 RFC 2464 [S], and the implications for MPEG-2 networks is described in draft-ietf-ipdvb-ar-xx.txt [Q].

A.4 Alternative methods of SI-SAP address scoping

An alternative method of SI-SAP address scoping is to use multiple independent instances of the SI-SAP to identify SI-SAP primitives that are intended for different "private" networks. This method can be used as well as, or instead of, the OUI-based methods.

Each SI-SAP would be used for traffic to/from a different "private" network (different groups of users). The impact of this method can be very localized in the protocol stack, since these multiple instances are only required to differentiate traffic at the SI-SAP and hence can be localized to the SIAF and SDAF. Above the SI-SAP the differentiation can be achieved at the IP layer using logically separate routers and physically separate networks, and below the SI-SAP this differentiation can be achieved by using different MAC logical channels, or similar lower layer features.

A typical ST is only required to support a single SI-SAP based on the assumption that a given ST is only required to support a single attached private network. The SI-SAP address scoping can therefore be defined via management configuration; e.g. by configuring each ST with the appropriate satellite dependent filtering of MAC channels. See clause A.5 for some specific examples.

However a Gateway ST (e.g. the hub ST in a star network) will be required to support multiple SI-SAPs. In the limit a single hub ST could be required to handle traffic with the full range of different scopes (i.e. traffic from all of the different networks) and hence support a separate (logical) instance of the SI-SAP for each of those networks.

NOTE: The same SI-SAPs can also be used to provide scoping/separation of the unicast traffic (to different networks) as well as scoping/separation of the multicast traffic.

A.5 Example of multicast address scoping

A.5.1 Multicast address scoping in DVB-RCS networks

As an example of a possible implementation of multiple OUIs by a single satellite network, we consider the case of DVB-RCS networks.

In this case the separation of the various scopes could be implemented by using several PIDs (or even ISI values for S2) where each PID contains multicast traffic for just one multicast OUI. Each OUI group can use the standard IANA (RFC 1112 [3]) mapping of individual multicast addresses, with the separation being achieved by the use of a "private" PID. Each ST then uses the standard PID filters (below the MAC level) to separate the wanted multicast flows: these PID filters are a normal feature in all DVB-RCS receivers.

NOTE: This approach resembles that of 3GPP and DVB-H.

Annex B (informative): Description of SI-SAP Primitives

B.1 C-Plane Group Receive Primitives

The Group Receive Open and Close primitives are used by the SI upper layers to enable or disable the reception of multicast groups in the lower layers, by supplying a list of the requested BSM_GIDs.

These primitives are only required in the case where the lower layers operate in a non-promiscuous mode and therefore explicit enabling or multicast group reception is required.

NOTE 1: The use of separate Open and Close primitives is based on the functional model defined in TS 102 357 [2]. This use of separate open and close primitives is designed to provide a simple mapping between the corresponding "include" and "exclude" lists in the IETF defined IP group management messages.

The Group Receive Status primitive is used to enable the upper SI-layers to request a list of all the currently enabled BSM_GIDs.

NOTE 2: The Status primitive is a new (proposed) primitive that is not identified in TS 102 357 [2] (V1.1.1).

The Group Receive primitives should only be used at the Egress STs.

B.2 Primitive definitions

The suffix on the primitive names in these tables refers to one or more of the primitive types, namely REQUEST (-req), CONFIRM (-cfm), INDICATION (-ind) or RESPONSE (-res). Each service uses one or more different types of primitives: the primitive types that are used for each service are explained in the following three clauses.

The presence or absence of each parameter in each type of primitive is defined using the following codes:

- Mandatory (M). A Mandatory parameter shall appear in all instances of that primitive.
- Optional (O). An Optional parameter may appear in a given instance.
- Equal (=). An equal symbol means that this parameter shall have the same value as the invoking parameter (i.e. it is a copy of the invoking parameter). A parameter in a Request primitive cannot have this status. This additional symbol can apply to either Mandatory (M=) or Optional (O=) parameters.
- Absent (-). An absent parameter shall not appear in that primitive.

B.2.1 SI-C-RGROUP_OPEN

The SI layer can enable a list of one or more GIDs at any time using the SI-C-RGROUP_OPEN-REQ primitive, containing an unordered list of the requested BSM_GIDs.

The SD layer should respond to the request using the SI-C-RGROUP_OPEN-CFM primitive which contains a copy of the list of BSM_GID together with a set of updated Cause Codes to indicate if each GID was successfully enabled.

Table B.2.1: SI-C-RGROUP_OPEN primitives

PRIMITIVE NAME	SI-C-RGROUP_OPEN-***
FUNCTION	Enable reception of the listed multicast groups

Primitive parameters:	**-req**	**-cfm**			**Comments**
RGROUP Query Handle	M	M=			Used to match request and confirm
Number of GIDs [N]	M	M=			Indicated the total number of GIDs
BSM_GID#1	M	M=			First BSM_GID; unordered list
Cause code#1	M	M			Success/ Failure indication Gives cause in case of failure
BSM_GID#2	M	M=			Second BSM_GID; unordered list
Cause code#2	M	M			Success/ Failure indication Gives cause in case of failure
../..					
BSM_GID#N	M	M=			Last BSM_GID; unordered list
Cause code#N	M	M			Success/ Failure indication Gives cause in case of failure
M					

B.2.2 SI-C-RGROUP_CLOSE

The SI layer can disable a list of one or more GIDs at any time using the SI-C-RGROUP_CLOSE-REQ primitive, containing an unordered list of the requested BSM_GIDs.

The SD layer should respond to the request using the SI-C-RGROUP_CLOSE-CFM primitive which contains a copy of the list of BSM_GID together with a set of updated Cause Codes to indicate if each GID was successfully disabled.

Table B.2.2: SI-C-RGROUP_CLOSE primitives

PRIMITIVE NAME	SI-C-RGROUP_CLOSE-*** (for the meaning of the wildcard "***" refer to annex A of TS 102 357 [2])
FUNCTION	Disable reception of the listed multicast groups

Primitive parameters:	**-req**	**-cfm**			**Comments**
RGROUP Query Handle	M	M=			Used to match request and confirm
Number of GIDs [N]	M	M=			Indicated the total number of GIDs
BSM_GID#1	M	M=			First BSM_GID; unordered list
Cause code#1	M	M			Success/ Failure indication Gives cause in case of failure
BSM_GID#2	M	M=			Second BSM_GID; unordered list
Cause code#2	M	M			Success/ Failure indication Gives cause in case of failure
../..					
BSM_GID#N	M	M=			Last BSM_GID; unordered list
Cause code#N	M	M			Success/ Failure indication Gives cause in case of failure

B.2.3 SI-C-RGROUP_STATUS

The SI layer can request a list of all currently enabled receive groups at any time using the SI-C-RGROUP_STATUS-REQ primitive.

The SD layer should respond to the request using the SI-C-RGROUP_STATUS-CFM primitive which contains a complete list of all BSM_GID that are currently enabled.

Table B.2.3: SI-C-RGROUP_STATUS primitives

PRIMITIVE NAME	SI-C-RGROUP_STATUS-***
FUNCTION	Request a list of all currently enabled multicast groups

Primitive parameters:	**-req**	**-cfm**			**Comments**
RGROUP Query Handle	M	M=			Used to match request and confirm
Number of GIDs [N]		M			Indicated the total number of GIDs
BSM_GID#1		M			First BSM_GID; unordered list
BSM_GID#2		M			Second BSM_GID; unordered list
../..					
BSM_GID#N		M			Last BSM_GID; unordered list

B.3 Parameters

B.3.1 RGROUP Query Handle

A locally unique label (unique only in the context of the associated SI-SAP) that is used to associate a given RGROUP request primitive with the corresponding confirmation primitive. This label is assigned by the upper layers, and is copied by the lower layer into the associated reply.

B.3.2 Number of GIDs

A16 bit integer listing to total number of GIDs in this primitive.

NOTE: A 16-bit number is recommended to correspond to the maximum number of records permitted in IGMPv3 [4].

B.3.3 BSM_GID

The BSM_GID is the SI-SAP address that identifies a specific multicast group.

The BSM_GID is a 48 bit number as defined in the TS 102 357 [2].

B.3.4 Cause code

The cause code parameter is inserted between each BSM_GID in the CFM primitive to indicate if each GID was successful, or to give the cause in case of failure.

NOTE: This parameter is included in the request primitives containing the dummy "Request" value. This is done to maintain the same primitive structure in request and confirm.

Parameter	Value	Comment
Type; Integer (1 byte)	0 = Success	Used in Req primitive
	1 = Fail: unknown cause 2 = Fail: too many GIDs enabled 3 = Fail: unknown GID requested 4 = Fail: invalid GID requested 255 = Request (dummy value)	

B.4 Procedures

B.4.1 Normal operation

The SI layer shall set the Cause Code parameter to the default value of « Request » in the request primitive and the SD shall update this parameter in the confirm primitive to indicate success or failure.

In case of a failure to enable or disable the requested GID, the reply should indicate the cause of failure using the Cause Code parameter. In the case of multiple causes of failure, the relevant cause code in the response should indicate the highest numbered cause. If no specific cause can be identified it should indicate « Fail: unknown cause ».

The list of GIDs is an unordered list and the lower layers may process the list in any order. However the complete list of GIDs shall be processed before any reply is given and the confirm primitive shall contain a valid response for every GID.

B.4.2 Exception handling

Any duplicate request (i.e. a request to enable or disable a GID that is already enabled or disabled) shall be treated as a valid request, and the confirm primitive for that GID should return the appropriate applicable cause code (e.g. « Success » error code in the case that the wanted outcome is already achieved).

The handling of conflicting requests is lower layer dependent. In the case of limited resources a later request may result in the invisible disabling of a previously enabled group (e.g. to free up resources for the newly enabled groups). Equally lower layer filtering limitations may result in additional GIDs being enabled in addition to those requested. In all such cases, the lower layers are not required to give any indication of any such unsolicited changes to the GIDs to the upper layers, other than responding to the present request by providing a valid Cause Code response for every GID.

Annex C (informative): ABBI Multicast System

C.1 Introduction

This annex describes an industry sponsored system called ABBI in the present document. It is a predecessor to DVB-RCS. ABBI Multicast is by design a star-push system (see clause 4.2). The main components of the ABBI multicast system comprise:

- Ingress point of HUB and restricted to encapsulated unicast multicast in ST.
- Egress point consisting of only STs.
- Multicast Access Control.
- Multicast Control Management (IETF messages).
- the Satellite Terminal (ST).

The Hub in ABBI is the only entity allowed to transmit true multicast through the forward link system.

C.1.1 Ingress Node

Multicast traffic arrives at the ABBI System two distinct ways. Either it is tunnelled from a remote source or it arrives as native multicast packets at the HUB. The other why is to send multicast encapsulated (in unicast) from the ST to the Multicast Server. The Hub processes IGMP reports from the IP-DVB GW out the Multicast Enabled Router to start a multicast flow in the ABBI system from the Internet. The NCC will use access control to set up the required lower layer multicast addresses in the IP-DVB Gateway and the MMT generating device that signals the IP/MAC distribution trees to STs.

C.1.2 Egress Nodes

The Egress nodes of the ABBI System are STs only, and all receive the same multicast flow. BSM Access Control is the same for all STs; if a multicast entry exists in the Multicast Mapping Table (MMT) then any Egress ST is allowed to replicate the multicast traffic on the premises network.

Egress STs may processes IGMP or PIM messages locally to control the forwarding of multicast groups. This implies that STs can have dynamic group membership but only to the groups to which the Hub is already subscribed, and this membership has no impact on the traffic forwarded over the BSM.

C.2 System Description

C.2.1 Multicast Access Control (MACD)

The Multicast Access Control is responsible for all HUB devices, including multicast transmission. The MACD has a GUI to start, stop, and schedule multicast flows; The MACD has the PID and address for each multicast service to be played out. It is the main component for adding, modify, or removing multicast services from the system.

C.2.2 Multicast Address Management

C.2.2.1 Multicast Mapping Table (MMT)

The MMT is table transmitted in the forward link from a simple DVB-SI ASI table play out device into the IP-DVB Encapsulator with information of information for each multicast IP address and PID. This defines the distribution tree for the ST to lookup the PID for a given multicast group address. This is employed since the Egress ST is not allowed to modify the distribution tree in the BSM Network as described in a star-push model.

C.3 Multicast Control Management

C.3.1 Multicast Enabled Router in HUB

Consumer Multicast Traffic is to be sent to the ABBI users for consumption as a product. All native multicast traffic will be sent out through an IGMP enabled interface on the main router to the IP-DVB Gateway. All Tunnelled Multicast traffic will use standard GRE (RFC 1701 [Z] and RFC 1702 [AA]) or IP-in-IP (IP protocol number 4) from the ISP location to the GRE or IP-in-IP decapsulator. The extracted multicast messages would then be forwarded to the IP-DVB Gateway via the firewall. The Hub also provides a Multicast server, which is maintained under Firewall protection.

C.3.2 IP-DVB Gateway

The IP-DVB Gateway (IPE) will distribute the multicast in the forward link as defined by the MMT. The address translation table in the IP-DVB Gateway is populated with multicast addresses and PIDs by the MACD. When a multicast packet arrives at the IPE traffic interface, the address is matched to a PID and service. It is then transmitted according to matching multicast service (including QoS values). Multicast traffic should be forwarded to all STs at this point.

The Egress ST reads the MMT and creates an address translation table. It receives and processes IGMP reports from the attached customer premises network. It will blindly forward the multicast group to the network if a match is made it is translation table. It will generate periodic IGMP queries to determine if a group member is still in the network.

C.4 Multicast Sessions Operation

Multicast session operation has two components in Phase 1: a ST component and a Hub component.

C.4.1 Invoking a Multicast Session

C.4.1.1 Egress ST-side

When a user wishes to join a multicast session (Multicast session means PID/IP pair), the user starts an application utilizing IP multicast. The Host application uses the Class D address and sends an IGMP JOIN message over the local (Host's) network. The Egress ST picks up the JOIN message and uses the IP address to look up the corresponding PID in its MMT table. When a match is made, the PID is put into the Egress ST's active MMT PID list so that the Egress ST can begin to decode the multicast data when it is broadcast from the Hub. If the number of Egress ST-supported multicast sessions is exceeded (minimum 14 multicast IP addresses, which can be spread over the maximum of 10 different Multicast PIDs per Egress ST assigned to multicast, no multiplex of IP multicast addresses over MAC multicast addresses is assumed on the same PID), the Egress ST ignores the JOIN message received from the host and no error message is given to the user by the SIT. All error processing for such a case by the application.

This is depicted in Figure C.4.1.

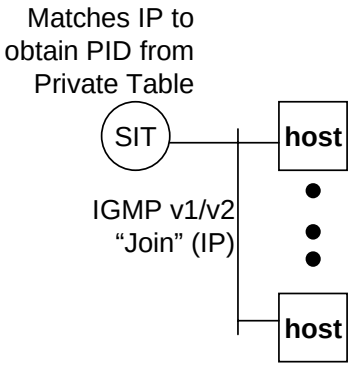

Figure C.4.1: Invoking a multicast receive session: SIT side

The Egress ST calculates the Multicast MAC address based on RFC 1112 [3], section 6.4. The ST parses the PMT where it will get the PID value carrying the Multicast Mapping Table (MMT) identified by the Elementary Loop with the stream type = 0x05 and table id = 0xC0 in the RCS content descriptor. From the MMT the Egress ST will know on which PID(s) it will get its multicast IP traffic.

The Egress ST maintains a counter of all active hosts on a particular group. The counter is incremented with each JOIN, and decrements with every LEAVE, with the PID being retired from the active list when the counter reaches zero.

C.4.1.2 Hub-side

When it is time to start a multicast session at the Hub (i.e., when transmission must begin), the ABBI operator goes to the MACD Table, highlights or double-clicks on the IP address for the multicast session to be started and selects the "start multicast" function (e.g., menu option).

Upon operator invocation of the start multicast function, the MACD performs the following tasks:

- Calculates the MAC addressed based on RFC 1112 [3], section 6.4 for inclusion in the SNMP Set message referred to in the step below.
- Uses the IP address and the derived MAC address to select the IP-MAC-PID combination and sends an SNMP message to the IP/DVB Gateway with the IP-MAC-PID triple. The IP-DVB Gateway detects a Class-D address, adds the triple to the Gateway table and sends a JOIN message to the router. Also, the Gateway pretends to be a host; that is, it responds to periodic polls by the multicast router about which hosts remain on what groups.

Upon receiving a JOIN message, the router propagates the JOIN message to other network entities. This is illustrated in Figure C.4.2.

Periodically, the IP-DVB Gateway resends the MPEG Multicast Mapping Table based on the time frequency specified by the parameter at the MACD. The cycling frequency is a configuration parameter that, like other configuration parameters, cannot be changed during IP-DVB Gateway operation.

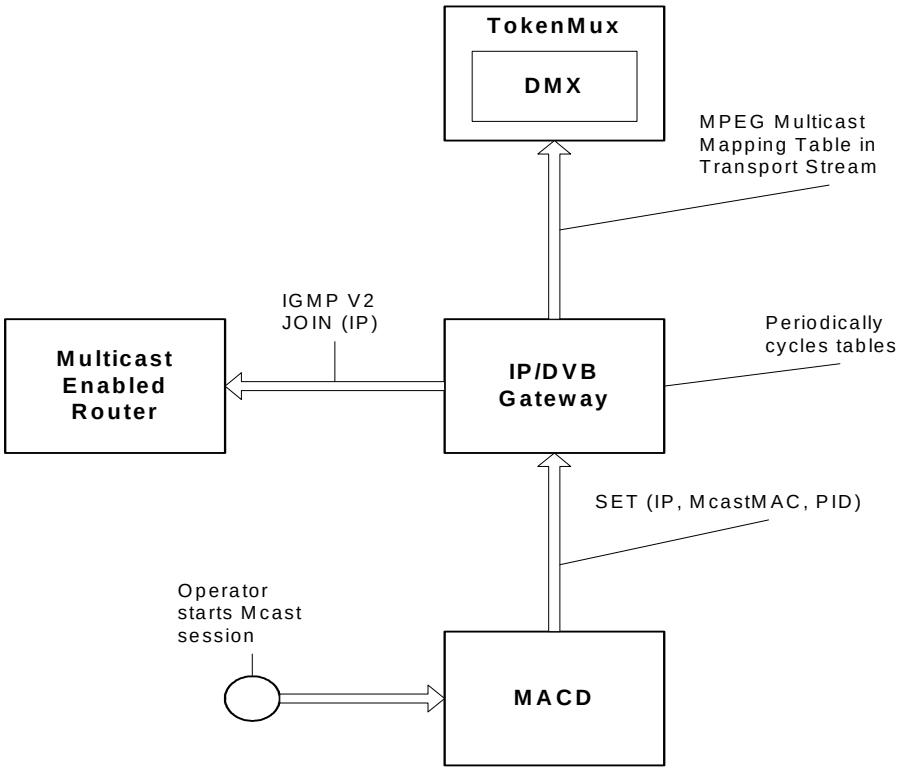

Figure C.4.2: Invoking a multicast receive session: Hub side

C.4.2 Revoking a Multicast Session

C.4.2.1 Egress ST-side

When a user wishes to leave a multicast session, the user stops the application using the IP multicast session. The Host application uses the Class D address and sends an IGMP LEAVE message over the local (Host's) network.

The Egress ST sends queries for hosts on every active group on a periodic basis. If there are no more listeners, the Egress ST can stop filtering on that multicast MAC address. If there is no more MAC filtering on this PID, the Egress ST removes the PID from its active cache.

C.4.2.2 Hub-side

When it is time to finish a multicast session at the Hub (i.e., when transmission must end), the ABBI operator goes to the MACD Tables and highlights the IP address to be revoked then selects the "stop multicast" function (e.g. Menu option). Upon operator invocation of the stop multicast function, the MACD performs the following tasks:

- Calculates the MAC addressed based on RFC 1112 [3], section 6.4 for inclusion in the SNMP Set message referred to in the step below.
- Uses the IP address and the derived MAC address to select the IP-MAC-PID combination and sends an SNMP message to the IP-DVB Gateway with the IP-MAC-PID triple.

The IP-DVB Gateway detects the Class-D address and (1) removes the tuple and (2) sends a LEAVE message to the router.

Upon receiving a LEAVE message, the router propagates the LEAVE message to other network entities.

C.5 Multicast Source Transmission

Multicast sessions may be sourced from the following entities using the following formats.

a) From an Ingress ST to the Hub, using unicast encapsulation. The Ingress ST will block native multicast transmission from the hosts to the hub (since the SIT will never transmit any IP traffic with a class D destination address to the hub).

b) Via the Internet to the Hub, using unicast encapsulation.

c) Via a dedicated link into the Mrouter, using a native multicast or using unicast encapsulation.

d) From a multicast server in the ABBI Hub.

All encapsulated multicast transmissions go through a GRE Decapsulator which de-capsulates the packets and re-routes into the mrouter using native multicast, as shown in Figure C.5.1. GRE encapsulation is used for Multicast. GRE is described in RFC 1701 [Z] and RFC 1702 [AA].

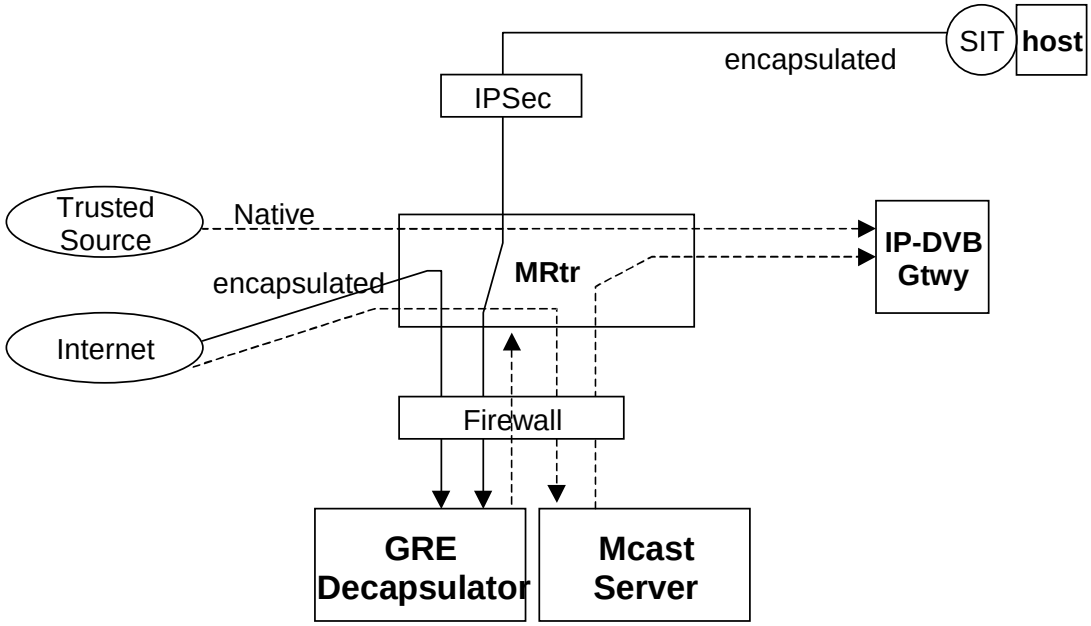

Figure C.5.1: Multicast source transmission model

The GRE Decapsulator receives encapsulated multicast transmissions. The Decapsulator strips the unicast encapsulation and forwards the multicast addressed packets through the firewall to the multicast router. The router then forwards the packets to the IP-DVB gateway (via the Ethernet switch). The GRE Decapsulator is used for streaming applications

The multicast server is the source for native local multicast flows and is used for store & forward applications. A customer should be able to submit a file to the server through the firewall from external networks. This is SES-supplied equipment.

Annex D (informative): RSM-A System

D.1 Multicast Service over RSM-A

RSM-A multicasting provides host user applications with (Internet Protocol) IP layer and multicast-capable interface access to the multicast/broadcast capability of the RSM-A satellite network.

D.1.1 Multicast Group Addressing

RSM-A multicasting, like IP multicasting, achieves efficient multipoint delivery through the use of Satellite Terminal (ST) virtual port groups, known in RSM-A as multicast groups. A multicast group is a group of one or more ST Virtual Ports (VPs) that is identified by a single Multicast Group Identification Number (MGID) destination address. An MGID has 000 in the three high-order bits (*Mcst bit-field)* of the *Destination ST Subaddress* in the ST identification field. The MGID is unique system-wide per satellite. Please see the RSM-A air interface specification for the format of the MGID and how it is encoded in the RSM-A packet header of a multicast RSM-A packet. A virtual port is equivalent to a Satellite Independent - Service Access Point (SI-SAP) and is identified by a unique MAC address or BSM ID. Virtual port and SI-SAP are used interchangeably in the present document.

D.1.2 Dynamic Group Management Signalling

As RSM-A is meant to be an IP-capable network and STs are required to seamlessly interface with IP hosts and routers, IP multicast protocols have to be supported at the terrestrial interface of STs. STs may use IGMP to learn multicast group membership. Dissemination of multicast routing information over RSM-A is performed through use of a multicast signalling protocol specifically designed for this purpose for RSM-A. This group management signalling can be used by connection-oriented multicast services to manage dynamic memberships for multicast sessions.

D.1.3 Multicast Services

RSM-A provides two broad categories of multicast services:

- Connection-oriented multicast services require the setup of a RSM-A connection in the system before any data can be distributed via multicast. Connection-oriented multicast is well-suited for multicast applications with stringent quality of service requirements generally characteristic of real-time multimedia applications.

- Connectionless multicast services do not require the setup of a connection in the system and deliver multicast data using connectionless delivery services. Connectionless services are best-suited for applications with a high tolerance for latency and jitter, such as non-real time data transfer applications. Connectionless services are not discussed in this paper.

The prerequisite for connection-oriented multicast communication is a Connection Setup process between the source ST and the NOCC.

The RSM-A Broadband Network system provides two types of connection-oriented multicast services. The services are:

- Scheduled Multicast Service - It is a multicast service with a scheduled start time in which a single source ST and one or more pre-configured ST receivers (i.e. static multicast groups) or one or more potential ST receivers (i.e. dynamic multicast groups) are pre-configured for a static or a dynamic multicast session, respectively, with the multicast session parameters. The configuration parameters include the session start time, prior to the start of the multicast session.

- On-demand Multicast Service - It is a multicast service with unscheduled start time in which multiple potential source STs and one or more pre-configured ST receivers or one or more potential ST receivers are pre-configured for a static or a dynamic multicast session, respectively, with the multicast session parameters. The configuration parameters include the start of session trigger classification, prior to the start of the session.

D.2 PIM-SM Support

PIM-capable STs (in short, PIM STs) appear on their terrestrial interface as a regular PIM-SM router. Native PIM-SM signalling is not, however, used on the RSM-A air interface. PIM STs perform the protocol inter-working require to translate PIM-SM control signalling to RSM-A multicast group management signalling, and vice versa.

For reasons of avoiding unwarranted complexity of PIM-SM implementation, in conjunction with reasons of operational efficiency, certain PIM-SM functions that are normally a part of PIM-SM implementation in IP routers are not needed in PIM STs. This clause identifies all such PIM-SM functions, and specifies the modified PIM functionality that PIM STs implement, with details being presented in a following clause. This later clause describes in detail the interactions between PIM-SM signalling and RSM-A multicast group management signalling. The approach taken therein is to describe all resultant changes from the behaviour that a PIM router would exhibit, rather than describing the complete set of multicast functions that have to be implemented. This helps to avoid duplication of material and also, the author believes, is a better way of presenting this design since the alternative would lead to a muddled distinction between the standard specification and RSM-A-induced artifacts of the design.

D.2.1 Location of Rendezvous Point

This design precludes the use of PIM STs as Rendezvous Points (RPs). PIM STs will not be candidate RPs for any multicast group. However, a solution is provided for the cases where a PIM ST might be directly attached to a multicast source, and no other router being present on the same network. The design allows for configuration of a multicast group with a *faux* or *virtual* RP. A source PIM ST can be instructed via configuration to effectively disregard a configured (the faux) RP and serve the multicast source in the absence of a real RP.

D.2.2 PIM-SM Bootstrap and RP-Information Distribution

This design precludes the use of a PIM ST as a PIM BSR (Bootstrap router). PIM STs will not be BSR candidates. However, PIM STs will process bootstrap messages.

All PIM STs (and PIM routers on their respective terrestrial LAN) are initially configured with the current RP information for each dynamic session. The NOCC is also configured with the current RP information for each dynamic session. Since RSM-A network administration requires that the BSR-set and the RP-set are on the same side of the RSM-A network as the source PIM ST, the source PIM ST detects when the RP changes. The content of the bootstrap 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. All PIM STs listen to BSR messages. When the source PIM ST detects a change in RP for a session, the source ST sends an RP Change Notification, in the form of a RSM-A management message, to the NOCC with the BSR message encapsulated in the management message. The NOCC transparently relays the encapsulated RP Change Notification as part of a management message and forwards it to the *Dynamic Multicast MGID* destination address using the Replication Group Number (RGN) of the multicast session.

Incidentally, the NOCC includes the RP information in Dynamic Join Acknowledgement/Announcement message that is sent to the *Dynamic Multicast MGID* destination address in the microcell when a Dynamic Join is received. The Dynamic Join Acknowledgement/Announcement message contains encapsulated BSR message if the RP for the group has changed. This allows PIM STs in the microcell to update their RP information and keep it current.

When a receiving PIM ST receives the announcement, it de-capsulates 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 source ST maintains a "RP Change" state per group, which is flagged when an RP changes. This allows the source 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 a receiving PIM ST, via the NOCC, may not match the RP for the specified group in its route entry.

D.2.3 Designated Routers and DR Election

The design allows for PIM STs to operate as DRs. PIM STs participate in DR election just as a normal PIM router would.

D.3 Scheduled Multicast Service

For Scheduled multicast service, RSM-A Broadband Satellite Multimedia System provides two types of the Scheduled multicast sessions. They are:

- Scheduled Static Multicast Session - It is a multicast session, in which the multicast group membership is pre-determined and does not fluctuate. Multicast group membership in scheduled static multicast service is limited to the pre-configured ST receivers. An ST **cannot** dynamically join a scheduled static multicast service.
- Scheduled Dynamic Multicast Session - It is a multicast session in which group membership can fluctuate (dynamic). An ST joins the multicast session as long as there is a host on its terrestrial network who wishes to receive datagrams and leaves the session when there is no host receiver on its terrestrial network.

D.3.1 Scheduled Multicast Connection Setup (SMCS)

Figure D.3.1 illustrates the Scheduled Multicast Connection Setup process.

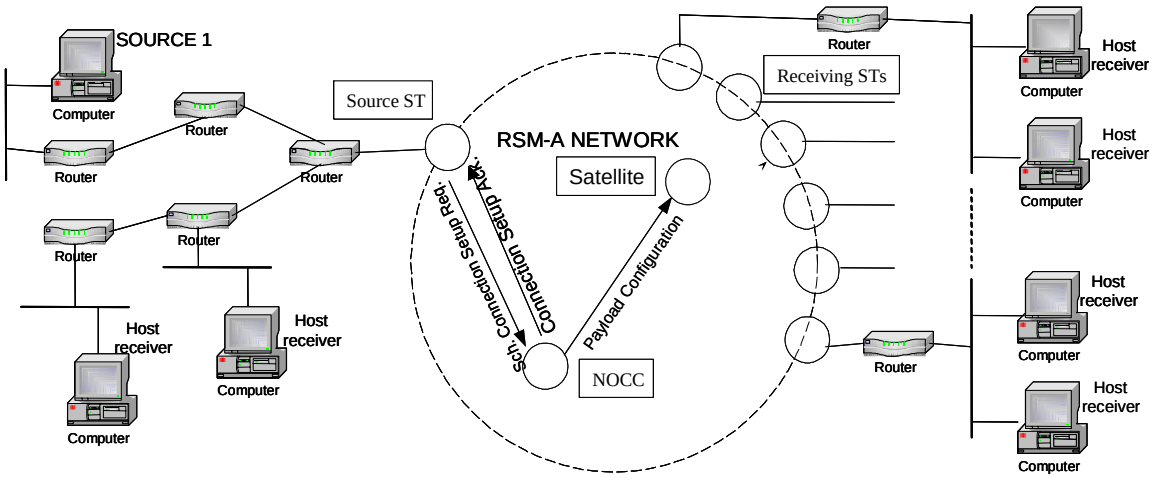

Figure D.3.1: Scheduled Multicast Connection Setup

A multicast-enabled ST is configured with the start time, port classification rules of the Scheduled multicast service, and the MGID destination address. The classification rules include a class D destination IP address. Scheduled multicast connection setup is always initiated by the source ST in a multicast session. When the Multicast Connection Setup request is received, the NOCC processes the request as follows:

- Verifies if the source ST is permitted to initiate a multicast connection request with associated parameters with admission control. Service may be denied based on lack of system resources, service access restriction or Community of interest restrictions.
- Resolves the multicast group members ST addresses to their destination downlink microcells.
- For static and dynamic multicast when ST receivers are present, it analyzes the microcell distribution for the group and determines the satellite broadcast replication network required for the multicast service.
- Configures the satellite payload, if necessary, according to the result of the analysis in the previous step. For Scheduled dynamic multicast session, the NOCC verifies if it has received a Dynamic Join from an ST in the multicast group for which the connection setup is being requested. If there is, at least, one receiving ST in the group who has joined, the NOCC proceeds with the processing of the connection setup request. If there is no active receiver, the NOCC sends a "Connection Setup In-progress" message to the source ST. The NOCC maintains a "Connection Setup In-progress" state for the group and checks the "Connection Setup In-progress" state with each Dynamic Join received. When a Dynamic Join from a receiving ST is received for the group, the NOCC deletes the "Connection Setup In-progress" state and resumes the processing of the connection setup request.

- Sends confirmation and session information to source ST after all the system, including security, and capacity checks are completed successfully.

After connection setup is successfully completed, the source ST enables the class D IP destination address, starts an end of session timer with and the multicast communication session begins. Figure D.3.2 illustrates multicast traffic flow for Scheduled Multicast Service.

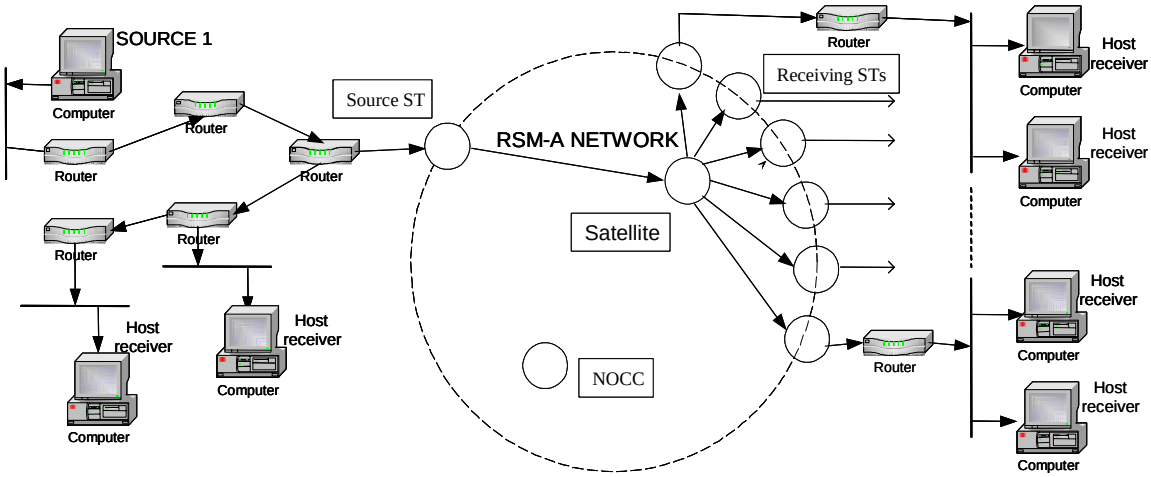

Figure D.3.2: Scheduled Multicast Traffic over RSM-A

A datagram, from the terrestrial network, arrives at the source ST from a host multicast server. The transport of the datagram from the multicast server to the source ST can be either static or dynamic.

If static, all routers and STs between the source (host sender of the multicast datagrams) and the host receivers must be statically configured to pass multicast datagrams. If dynamic, the source ST must join a multicast shared-delivery tree in order to receive the multicast datagram. After a successful connection setup, the source ST uses the pre-configured classification rules to map incoming user data to the scheduled multicast session.

The receiving STs, consisting of the Scheduled multicast group, are either statically configured or can dynamically request to join the session but not both. A Scheduled multicast session, with statically configured membership, does **not** require a multicast routing protocol to distribute the multicast datagram. At startup, or when a new scheduled multicast configuration update is received, receiving STs enable the MGID destination address to receive multicast datagrams. After confirmation of the Multicast Connection Setup request, the source ST starts sending datagrams. Upon receiving the datagrams, the receiving ST forwards them according to the configured active port list.

A Scheduled multicast session with dynamic destination ST membership may use IGMP to distribute group membership information as described in clause 6.1.

D.3.2 Forwarding Scheduled Multicast Datagram

The source ST monitors the destination address IP header field of all IP datagrams on each active user port. If a class D IP destination address is detected, it is compared with a list of configured class D IP destination addresses expected to arrive on the user port that the datagram is received. If the class D IP destination address of the datagram does not match any in the list for that user port, the datagram is silently discarded. If there is a match, the source ST maps the destination address to the corresponding Multicast Group Identification Number for that domain. Note that if processing has gotten this far, then successful connection setup for this scheduled multicast session has already occurred. The source ST constructs the appropriate RSM-A packet and forwards it towards the satellite with the MGID as the destination RSM-A address.

D.3.3 Receiving Scheduled Multicast Datagram

The receiving ST monitors the destination RSM-A address field in the RSM-A packet header of all RSM-A packets received on the air interface. If an MGID destination address is detected, it is compared with a list of configured MGID destination addresses expected to arrive on the air interface. If the MGID destination address of the packet does not match any in the list for the air interface, the packet is silently discarded. If there is a match and the session is static, the receiving ST forwards the datagram according to the output ST user port list configured in the multicast session information table. If the multicast session information indicates more than one output ST user port, the ST replicates the datagram and forwards it on each port in the output port list. Note that the ST assumes that the datagram is valid on the user ports on which it is forwarded since the configuration of the output list is provided by the NOCC. The NOCC deals with End User Group (EUG) connectivity restrictions. The ST does not. If there is a match and the session is dynamic, the receiving ST forwards the datagram according to the output ST port list, in the dynamically created route entry, indicating output ports with group members.

D.3.4 Disconnecting a Scheduled Multicast Service

After successful Scheduled Multicast Connection setup, the source ST starts an End of Service timer. When the timer expires, the source ST stops sending datagrams and sends a Scheduled Multicast Service Release message to the NOCC. The NOCC de-allocates all resources that it has allocated to the session and sends an acknowledgment to the source ST. Upon receiving an acknowledgment from the NOCC, the source ST frees all resources allocated to the session.

D.4 On-demand Multicast Service

For On-demand multicast service, RSM-A Broadband Network System provides two types of the On-demand multicast session. They are:

- On-demand Static Multicast Session - It is a multicast session, in which the multicast group membership is pre-determined and does not fluctuate.
- On-demand Dynamic Multicast Session - It is a multicast session in which group membership can fluctuate (dynamic). An ST, pre-configured as a potential member of a group, joins the multicast session as long as there is a host on its terrestrial network who wishes to receive datagrams and leaves the session when there is no host receiver on its terrestrial network.

D.4.1 On-demand Multicast Connection Setup

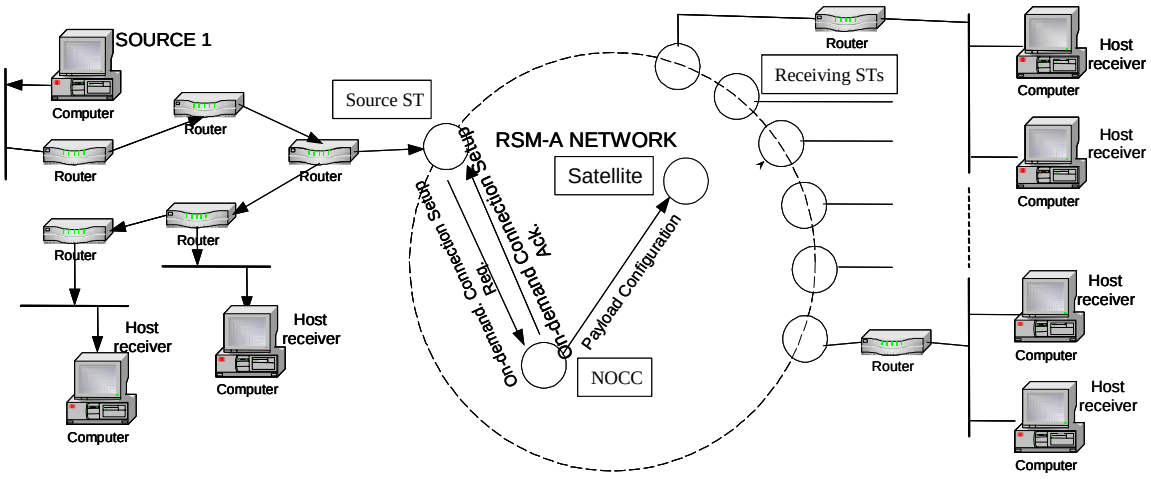

Figure D.4.1: On-demand Multicast Connection Setup

A multicast-enabled ST is configured with the classification rules of the On-demand multicast service and the MGID destination address. The classification rules include a class D IP destination address. On-demand multicast connection setup (OMCS) is always initiated by a potential source ST in a multicast session. On demand multicast connection setup may be triggered by data pattern in a specific IP multicast datagram or the class D IP destination address of the multicast datagram to setup a multicast connection. When the source of the trigger for multicast connection setup is detected in a datagram or one of the above triggers, and in addition for a dynamic session if a Dynamic Join Indication has been received from the NOCC, any of the potential source STs can initiate an OMCS by sending a Multicast Connection Setup Request, with connection parameters, to the NOCC using NOCC-ST management protocol. However, the NOCC admits the first connection request received and any subsequent connection requests are denied. The NOCC admits only one source ST per session. When the Multicast Connection Setup request is received, the NOCC processes the request as follows:

- Verifies if the source ST is permitted to initiate a multicast connection request with associated parameters with admission control. Determines if system capacity is exceeded based on the associated request parameters.
- For On-demand dynamic multicast session, the NOCC verifies if it has received a Dynamic 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 NOCC proceeds with the processing of the connection setup request. If there is no active receiver, the NOCC sends a "Connection Release" message to the source ST.
- Verifies the multicast connection capacity requirements of the request with system capacity. Service is denied if there is insufficient capacity.
- Resolves the multicast group members ST addresses to the destination downlink microcell.
- Analyzes the microcell distribution for the group and determines the satellite broadcast network required for the multicast service. If packet replication is required, the NOCC reserves a replication group number (RGN).
- Configures the payload, if required, according to the result of the analysis in the previous step.
- Grants request and sends confirmation and session information to source ST when all the system, including security, and capacity checks are completed successfully.

While the On-demand connection setup is in progress (i.e. the NOCC processes the request), the source ST discards multicast datagrams it receives for the group. After connection setup is completed, the multicast communication session begins. The source ST transmits multicast datagrams received over the air interface.

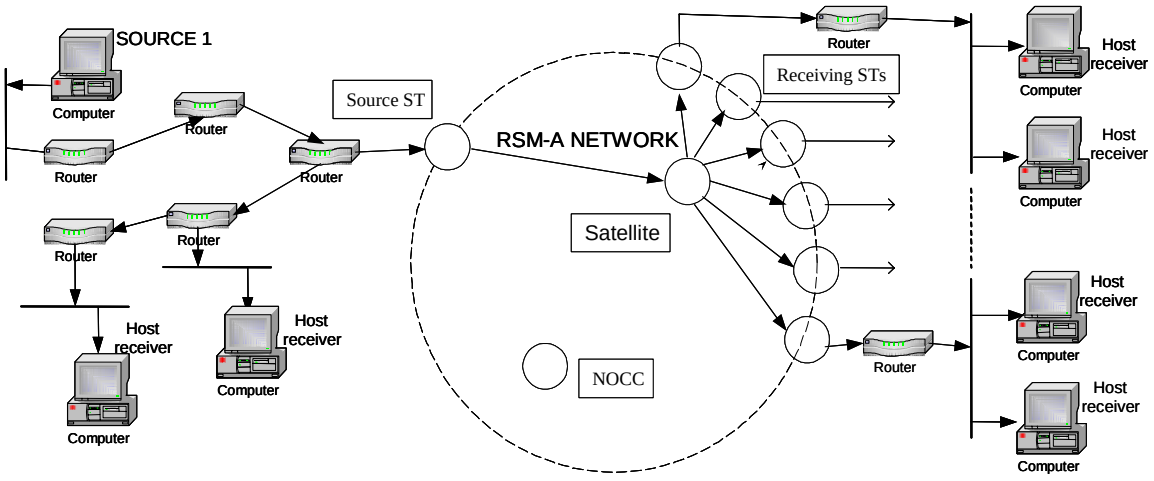

Figure D.4.2: On-demand Static Multicast Traffic over RSM-A

The receiving STs, consisting of the On-demand multicast group, are either statically configured or can dynamically request to join the session. An On-demand multicast session, with statically configured membership, does not require a multicast routing protocol to distribute the multicast datagram. At startup, or when a new scheduled multicast configuration update is received, receiving STs enable the MGID destination address to receive multicast datagrams. After confirmation of the Multicast Connection Setup request, the source ST starts sending datagrams. Upon receiving the datagrams, the receiving ST forwards them according to the configured eligible port list.

D.4.2 On-demand Multicast Datagram Forwarding

The source ST monitors the destination address IP header field of all IP datagrams on each active user port. If a class D IP destination address is detected, it is compared with a list of configured class D IP destination addresses expected to arrive on the user port that the datagram is received. If the class D IP destination address of the datagram does not much any in the list for that user port, the datagram is silently discarded. If there is a match, the source ST maps the destination address to the corresponding Multicast Group Identification Number (MGID) for that domain. Note that if processing has gotten this far, then successful connection setup for this On-demand multicast session has already occurred. The source ST constructs the appropriate RSM-A packet forwards it towards the satellite with the MGID as the destination RSM-A address.

D.4.3 Receiving On-demand Multicast Datagram

The receiving ST, in an On-demand multicast session, monitors the destination RSM-A address field in the RSM-A packet header of all RSM-A packets received on the air interface. If an MGID destination address is detected, it is compared with a list of configured MGID destination addresses expected to arrive on the air interface. If the MGID destination address of the packet does not much any in the list for the air interface, the packet is silently discarded. If there is a match and the session is static, the receiving ST forwards the datagram according to the output ST port list configured in the multicast session information table. If the multicast session information indicates more than one output ST user port, the ST replicates the datagram and forwards it on each port in the output port list. If there is a match and the session is dynamic, the receiving ST forwards the datagram according to the output ST port list, in the dynamically created route entry, indicating output user ports with group members.

D.4.4 Disconnecting an On-demand Multicast Service

After successful On-demand Multicast Connection setup, the source ST starts an end of session timer. The source ST restarts an On-demand Multicast Session Timer, for every datagram that it transmits, to timeout a session. The timer is restarted when a datagram is transmitted and cancelled when a datagram is received. When the timer expires, the source ST assumes that there are no more datagrams to be transmitted and sends an On-demand Multicast Service Disconnect message to the NOCC. The NOCC frees all resources that it allocated to the session and sends an acknowledgment to the source ST. Upon receiving acknowledgment from the NOCC, the source ST frees all resources allocated to the session.

D.4.5 RSM-A Multicast Group Management Signalling with PIM-SM

Convention: Throughout the present document the terms "input interface" and "output interface" have been used. The use of the terms "input interface" and "output interface", with reference to the source ST and the receiving ST, needs clarification. The naming convention is easily understood when referenced to the flow of multicast datagram traffic. Datagrams arriving on the input interface are considered upstream to the router. The datagrams forwarded on the output interface are considered downstream to the router. Multicast datagram traffic flows opposite in direction to PIM Join/Prune messages. For RSM-A multicast, it is recommended that the source of the multicast, the RP and the source PIM ST be on the same side of the satellite network to avoid multiple hops over the satellite network.

In general, the input interface points towards the Rendezvous Point (RP) and in some cases towards the source (host) of the multicast datagram. This is the upstream interface where PIM Join/Prune messages are forwarded towards the RP. This is also the interface on which multicast datagrams arrive. The output interfaces point towards neighbouring downstream routers. The output interface is where PIM Join/Prune messages arrive at the ST and it is the interface where multicast datagrams are forwarded (leave the ST) towards group members for reception.

Figure D.4.3 illustrates the flow of traffic for join messages and multicast datagrams and uses the flow to show the input and output interfaces.

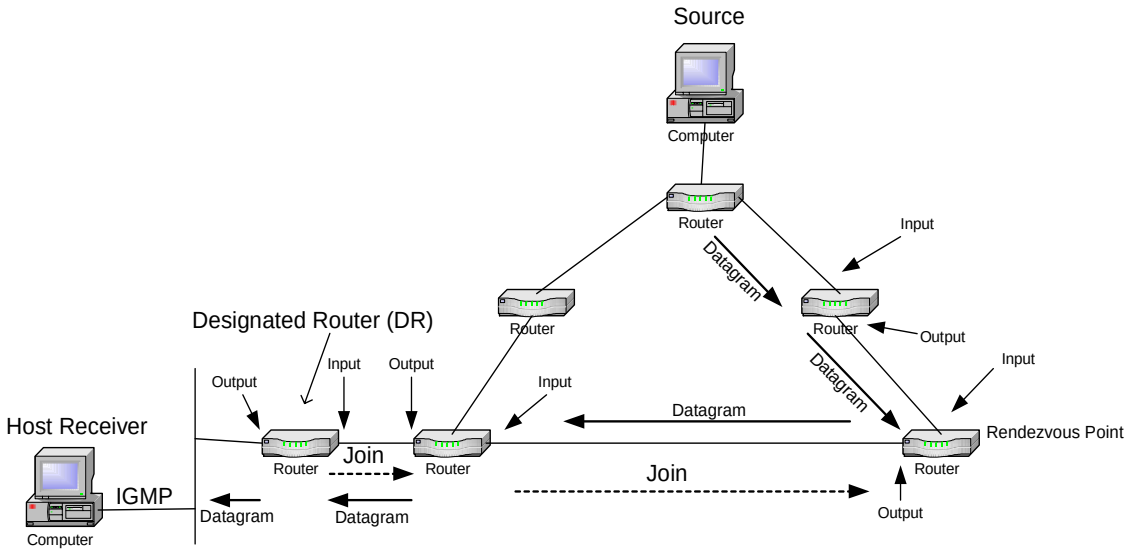

Figure D.4.3: PIM-SM Message Input/Output Convention

RSM-A Specific Considerations:

The input interface, for the source ST, is the ST user port. The source ST user port is where multicast datagrams arrive. The source ST user port points up towards a neighbouring upstream PIM router or the RP and it is where PIM Join/Prune messages are forwarded towards the RP. The input interface, for the receiving ST, is the RSM-A air interface. For the receiving ST, the RSM-A air interface is where multicast datagrams arrive and it is where (Dynamic) PIM Join/Prune messages are forwarded towards RP (via the NOCC).

The output interface, for the source ST, is the RSM-A air interface. For the source ST, the RSM-A air interface is where (Dynamic) PIM Join/Prune messages arrive (via the NOCC) and it is where multicast datagrams are forwarded (leave) towards group members for reception. The output interface, for the receiving ST, is the ST user port. For the receiving ST, the output interface is where multicast datagrams are forwarded (leave) towards group members for reception and it is where PIM Join/Prune messages arrive.

Note that RSM-A BSM system supports multicast routing purely over terrestrial interfaces (i.e. ST user port-to-ST user port without going over the air interface) therefore there could be multiple ST user ports functioning as input or output interfaces for *different* multicast sessions. Even then, configuration of all valid multicast groups at STs is a pre-requisite for the ST to participate in such purely terrestrial multicast routing.

D.4.5.1 Dynamic Join for a Multicast Group - Scheduled or On-demand

Figure D.4.4 illustrates RSM-A Dynamic Join and the propagation of IGMP group membership information and PIM-SM Join/Prune messages towards the RP. At the receiving PIM ST end of the network, there are two host groups (Group 1 and Group 2) on separate LANs connected to a PIM router, the designated router (DR). Each host computer generates and sends an IGMP Group Membership Report to the group address, which is received by the PIM router. Since the PIM router's unicast routing table indicates that the ST is the next hop towards the RPs, the PIM router generates and forwards a PIM-SM Join message upstream, on its input interface, to the PIM ST towards the RP. The PIM ST generates and forwards a Dynamic Join Request to the NOCC. The NOCC uses the MGID provided in the Dynamic Join Request to locate the source PIM ST and forwards a Dynamic Join Indication to the source PIM ST. (Note that the Dynamic Join Indication is only sent if the Dynamic Join Request is the first join request to be received for that multicast group. All subsequent Join Request messages are locally acknowledged by the NOCC.) The source PIM ST uses the information provided in the Dynamic Join Indication to generate and forward PIM-SM Join message to the RP. In this illustration, host GROUP 1 receives multicast datagrams from SOURCE 1 and host GROUP 2 receives multicast datagrams from SOURCE 2. Note that multicast data traffic flows in the opposite direction to PIM Join/Prune messages towards the DR.

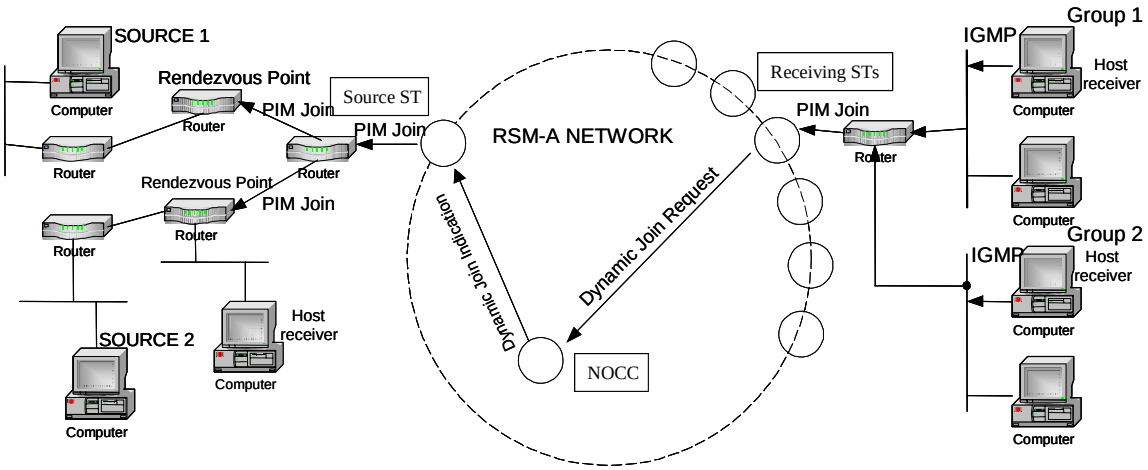

Figure D.4.4: RSM-A Dynamic Join

When a multicast-enabled PIM ST receives an IGMP group membership indication from a local host directly attached to its subnet or an explicit PIM Join message from a directly attached PIM router, for a group for which it has no entry, the PIM ST router maps the group address to one of the Candidate - Rendezvous Points (C-RPs). The entry for the C-RP in the routing table must include the group to which the C-RP is mapped. It is important to note that the PIM ST that receives the IGMP group membership indication or the PIM Join message is the *receiving* ST. However, under certain conditions the source ST may receive an IGMP message on one of its terrestrial ports and route PIM Join/Prune message to another port towards the RP. Thus, it is acting as a receiving ST. Subsequently, the receiving PIM ST generates a (RSM-A) Dynamic Join Request message and forwards the Dynamic Join Request to the NOCC. Note that the Dynamic Join Request message is a Dynamic Join/Prune Indication message with a NULL prune list. The Dynamic Join Request message includes the RSM-A management message header information, the microcell number of the receiving ST, etc. and an encapsulated PIM Join/prune message in the data field of the message. Essentially, the receiving ST generates a RSM-A (PIM-SM) Dynamic Join Request by encapsulating the PIM Join/Prune message in a RSM-A management message and forwards the message to the NOCC.

The receiving ST maintains a Dynamic Join Request Timer to guarantee acknowledgment (i.e. reception) of the message. The NOCC de-encapsulates the PIM Join/Prune message from the Dynamic Join/Prune Request message. The NOCC generates a RSM-A Dynamic Join Indication message, which includes an encapsulated PIM Join/Prune message and sends it to the source ST in a Scheduled multicast or potential source STs in an On-demand Multicast. The encapsulated PIM Join/Prune message is the PIM message received with the Dynamic Join/Prune Request message. The NOCC maintains a Dynamic Join Indication Timer to guarantee acknowledgment of the message.

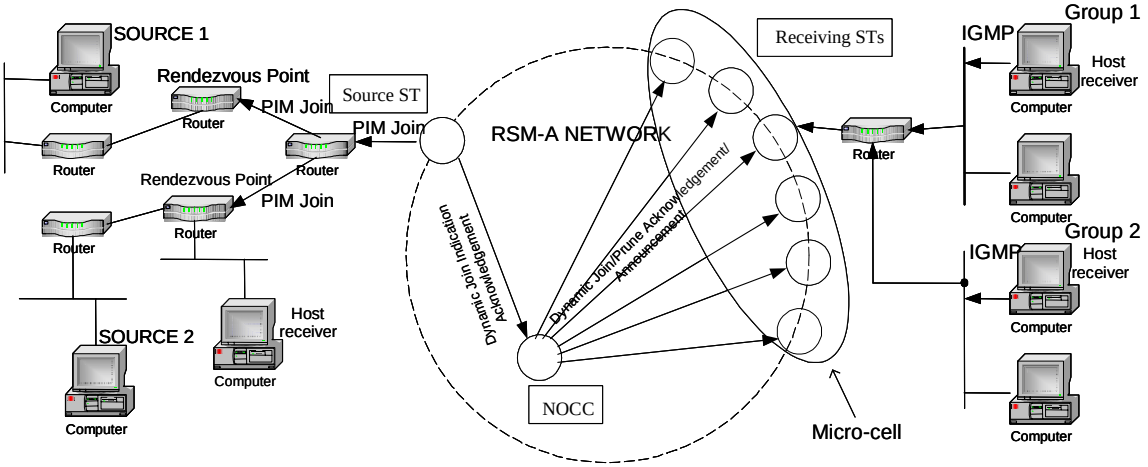

Figure D.4.5: RSM-A Dynamic Join/Prune Acknowledgement

In a Scheduled Multicast, the source PIM ST sends a Dynamic Join Indication Acknowledgement message to the NOCC to acknowledge the indication message. In an On-demand Multicast, each potential source PIM ST sends a Dynamic Join Indication Acknowledgement message to the NOCC, using a random delay timer, to acknowledge the Indication message. The NOCC processes the first Dynamic Join Indication Acknowledgement message received and discards the rest.

If necessary, the NOCC sends a configuration update to the Satellite Payload to update the replication table so that the receiving PIM ST can receive multicast datagrams. The NOCC microcasts back, to the microcell of the receiving ST, a Dynamic Join Acknowledgement/Announcement, per multicast group, to the group address, granting the Dynamic Join request. The NOCC microcasts the Dynamic Join Acknowledgement/Announcement message to the microcell of the receiving ST so that all the multicast group members, within that microcell, see it. Figure D.4.5 illustrates the NOCC-ST Dynamic Join/Prune Acknowledgement/Announcement process.

D.4.6 Dynamic Prune over RSM-A

PIM STs, with host group membership on directly attached networks, send IGMP Host/Group Membership Queries periodically to refresh their knowledge of memberships present on a particular network. The implementations of IGMPv1, IGMPv2 and IGMPv3 handle group members leaving a group differently, however. The implementations are explained as follows:

- If the PIM ST is operating in IGMPv1 mode and it does not receive IGMPv1 Host Membership Report for a particular group after a number of queries, the PIM ST assumes that the group has no local members.
- If the PIM ST is operating in IGMPv2 mode, the host that is the last host to send an IGMPv2 Group Membership Report on behalf of the group sends an IGMPv2 Leave Group message when it leaves a multicast group. The IGMP Leave Group message is sent to the *all-routers-group* multicast group address 224.0.0.2. When the message is received, the PIM ST sends an IGMP Group Specific Query to the multicast group address of the group being left and starts a timer. When the timer expires, the PIM ST assumes that the group has no local members if it does not receive an IGMP Group Membership Report.
- If the PIM ST is operating in IGMPv3 mode, a host, who wishes to leave a group, sends an IGMPv3 Membership Report message with the BLOCK_OLD_SOURCES record indicating the group and the source(s) that the host no longer wishes to receive from.

If the PIM ST has host group membership on a directly attached network, and a PIM neighbour on that network, and the PIM neighbour is the DR, the PIM ST receives a PIM Prune message. When the PIM ST no longer has local group members, it stops remote forwarding of multicast datagrams onto the local LAN for that group. The PIM ST receives PIM Prune messages from downstream PIM routers with no local members. A receiving PIM ST generates a Dynamic Prune Request and forwards it to the NOCC if and only if it has no output interface with members belonging to the group. Note that the Dynamic Prune Request message is a Dynamic Join/Prune Request message with a NULL join list. The Dynamic Prune Request message includes the RSM-A management message header information, the cell number of the receiving ST, etc. and an encapsulated PIM Prune message in the data field of the message. The receiving ST maintains a Dynamic Prune Request Timer to guarantee acknowledgment (i.e. reception) of the message.

When Dynamic Prune Request is received, the NOCC microcasts back, to the microcell of the receiving ST, a Dynamic Prune Acknowledgement/Announcement to the group address. The NOCC microcasts the Dynamic Prune Acknowledgement/Announcement message to the microcell of the receiving ST so that all the multicast group members, within that microcell, see it. Figure D.4.5 illustrates the NOCC-ST Dynamic Prune Acknowledgement/Announcement process.

When receiving PIM STs receive Dynamic prune Acknowledgement/Announcement for a group, the prune list is checked. If the list contains a source or RP for a group for which a Dynamic Prune message has been scheduled, the scheduled Dynamic Prune Request message is cancelled and the route entry is deleted for that group. If the prune list contains a source or RP for which the receiving PIM ST has a corresponding active route entry, then a Dynamic Join for the group is sent to the NOCC to override the prune.

The NOCC maintains group membership "state" by monitoring the number of microcells that have members of a group. When the NOCC receives a Dynamic Prune Request, it checks the group membership state. When a group membership drops to zero in a microcell, the NOCC sends a message to the payload to update the RGN. This stops datagrams, destined for a group, from being broadcast to a microcell with no group members. When a group membership, in all cells, drops to zero, the NOCC must either suspend or disconnect the service.

For On-demand Multicast session, the NOCC sends an On-demand Multicast Service Release Indication message, with the PIM Prune message encapsulated, to the source ST. The source ST stops transmitting datagrams to the group. The source generates a PIM Prune message and uses its unicast routing table information to forward it hop-by-hop towards the RP to prune the multicast shared delivery tree between it and the RP. For the special case where the source PIM ST is directly connected to the host of the multicast source (virtual Rendezvous Point), the PIM Prune message is not generated.

For Scheduled Multicast Session, the NOCC is configured whether to disconnect the session or suspend the session until a receiver joins if the multicast service duration has not expired. If configured to disconnect the session, the NOCC sends a Scheduled Multicast Service Release Indication to the source ST. The source generates a PIM Prune message and uses its unicast routing table information to forward it hop-by-hop towards the RP to prune the multicast shared delivery tree between it and the RP. For the special case where the source PIM ST is directly connected to the host of the multicast source (virtual Rendezvous Point), the PIM Prune message is not generated.

When the session ends, the source ST sends an acknowledgment back to the NOCC. The source ST ends the session and frees all resources allocated for the session. The NOCC frees all resources that it allocated to the session when it receives the acknowledgment.

D.4.6.1 Processing Dynamic Join Request in Response to a Dynamic Prune Announcement

When the NOCC receives a Dynamic Join Request message after broadcasting a Dynamic Prune Announcement, it processes the Dynamic Join Request message in the same manner as any Dynamic Join Request message it receives from a receiving ST as a result of a PIM Join or IGMP message. Since this is not the first join (i.e. group membership already exists), the NOCC does not forward a Dynamic Join Indication message to the source ST. However, the NOCC acknowledges the Dynamic Join Request by sending the acknowledgment/announcement message to the microcell of the receiving ST to the *DynamicMulticast MGID*.

The NOCC maintains a Dynamic Join Acknowledgement/Announcement Hold Timer, which indicates the amount of time the Dynamic Join Acknowledgement/Announcement message is valid. While the timer is running, any Dynamic Join Request received for the group by the NOCC is ignored. When the timer expires, the NOCC assumes that PIM ST group members that are active (up) and capable of receiving the message might have received it. After the timer expires, the NOCC starts processing subsequent Dynamic Join Request message received.

When the receiving STs receive Dynamic Join Acknowledgement for a group, any scheduled Dynamic Join Request message for the group is cancelled. Any receiving ST who is a potential receiver in the group and has not received a PIM Join message, creates a route entry for the group with NULL output so that multicast datagrams are not forwarded. In addition, any joins received for that group does not trigger a Dynamic Join Request to the NOCC but triggers the ST user port on which the PIM Join arrived to be added to the output port list. Note that, although this changes the output interface list from NULL to non-NULL it does not trigger a Dynamic Join Request to be sent to the NOCC.

D.5 PIM-SM Implementation

This clause provides details of implementation of PIM-SM in PIM STs, and on inter-working of PIM-SM signalling with RSM-A group management signalling.

D.5.1 Overview of Multicast Delivery Tree Establishment

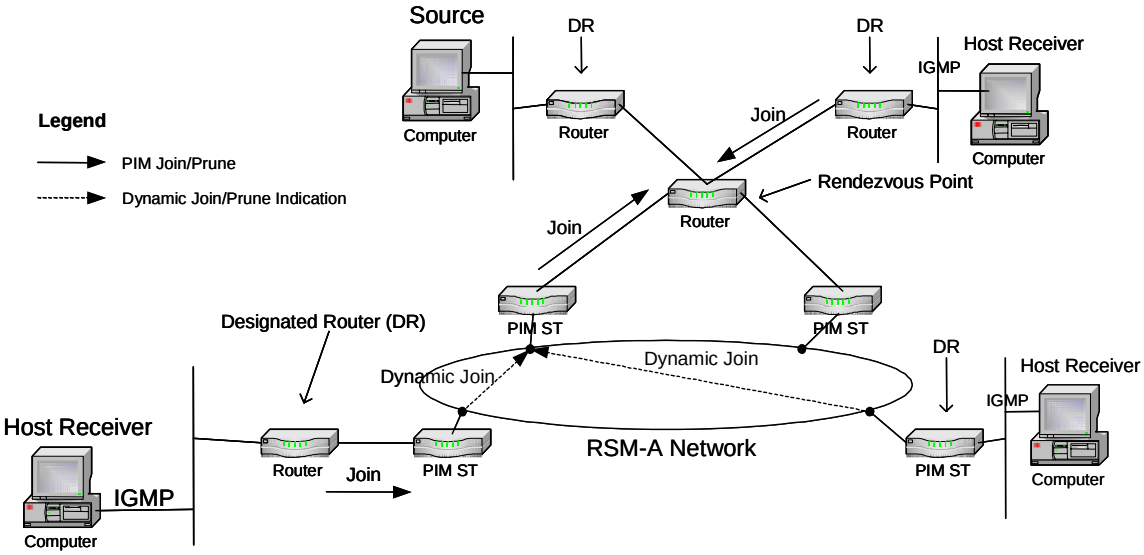

Figure D.5.1: Multicast Delivery Tree Establishment

Figure D.5.1 illustrates multicast delivery tree establishment in the context of RSM-A. The process is similar to delivery tree formation (via PIM-SM) in terrestrial IP networks, the difference being that over the RSM-A network instead of PIM-SM a RSM-A multicast group management protocol is used. Instead of native PIM-SM messages, PIM STs send Dynamic Join Indication message to the NOCC (not shown) with the PIM Join/Prune message encapsulated in the Dynamic Join Indication message. The NOCC processes the message and forwards a Dynamic Join Indication to the source PIM ST to complete a tree branch, keeping with the tree terminology.

Normally the DR performs a deterministic hash function over the sparse-mode region's current RP-set to uniquely determine the RP for this group. However, generally PIM STs will be configured with one RP and therefore it may not be necessary for the PIM STs to perform this function. However, it the hash function is performed, the group should map to the configured RP.

D.5.2 Multicast Datagram Processing and Forwarding

Figure D.5.2 illustrates the multicast data forwarding and routing process in RSM-A. It shows two multicast sources, SOURCE1 and SOURCE2, with associated rendezvous points, Rendezvous Point1 and Rendezvous Point2. Each source transmits the datagram on the subnet. The datagram is received by the corresponding DR who forwards it to its respective RP. At the RP, the datagram is distributed down the shared delivery tree. The source PIM ST is on the delivery tree and therefore receives the datagram. The source PIM ST transmits the datagram over RSM-A and two PIM STs who have members on their respective terrestrial network receives it.

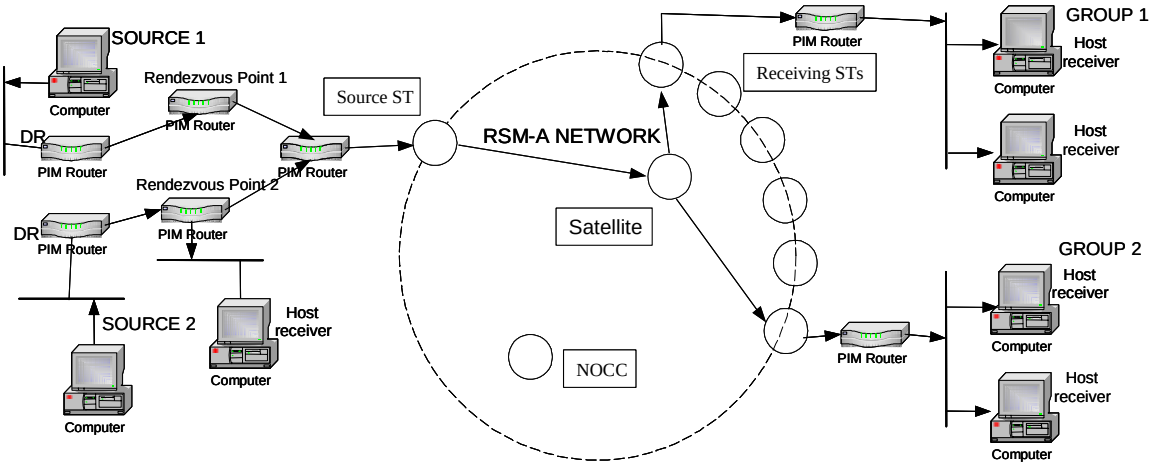

Figure D.5.2: Multicast Datagram Forwarding in RSM-A using PIM-SM STs

The processing of multicast data packets received on an interface is the same as that specified earlier for PIM routers

A source PIM ST, with directly connected host(s) may switch to the shortest path tree if it is configured to switch. However, it must first join the shared delivery RP-tree. Source PIM STs can switch to a source's shortest path tree (SP-tree) after receiving datagrams from that source over the shared RP-tree. The recommended policy is to initiate the switch to the SP-tree after receiving a significant number of datagrams during a specified time interval from a particular source. To realize this policy the source PIM ST monitors data packets from sources for which it has no source-specific multicast route entry and initiates such an entry when the data rate exceeds a pre-configured threshold.

Figure D.5.3 illustrates the switch to the shortest path tree process. The source PIM ST creates a (S, G) state first, and then switches to the shortest path tree (SP-tree) to the source multicast host, and removes itself from the shared delivery RP-tree. The source PIM ST switches to SP-tree if and only if it is configured to switch.

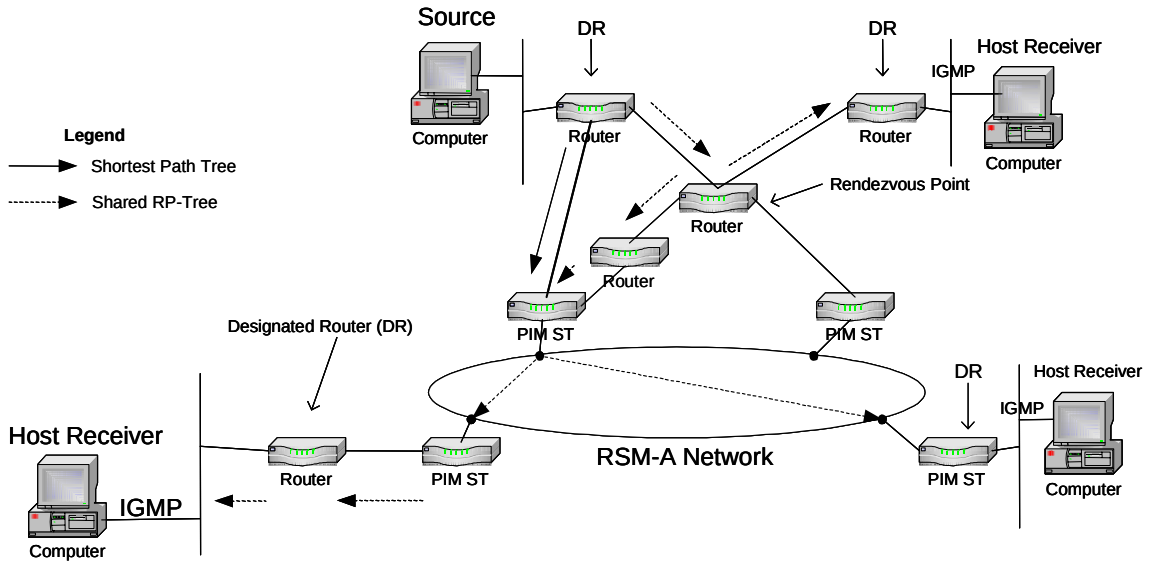

Figure D.5.3: PIM-SM - Switching to Shortest Path Tree

Annex E (informative): RSM-B Satellite Multicast System

E.1 Introduction

In a BSM network where the networks attached to the STs have multicast end nodes (in fact, this is the most usual case), the satellite network will not behave as an Internet multicast transit network. Rather, external multicast traffic will be injected through feeders/gateways. In these cases IGMP proxying is enough to provide multicast services to the attached networks.

In other cases the organization of networks attached to terminals may be more complex and involve also multicast routing protocols like the ones enumerated above. In some proposals even a convergence signalling protocol is defined between the Internet multicast routing protocol and the DVB signalling. But terrestrial networks attached to terminals finally behave as "leaf nodes" in the multicast tree, and a solution based on IGMP at the terminal can be used to signal group membership to the satellite enabled multicast network.

There are several issues to take into account when using IGMP in a satellite system.

- IGMP flooding issue.

The IGMP protocol is well suited for a terrestrial network with a shared medium. Indeed every host can listen to the reports transmitted by the others. If a report is already transmitted, the host will also stop the ongoing timer and will not send its report. This technique will prevent hosts from replying useless reports. However, it is not the case in a satellite context: satellite terminals can not directly listen to the replies from other member terminals. Moreover, as the satellite multicast groups can be very large (satellite networks with 5000 STs) and very dynamic, this issue can have some serious consequences: the IGMP querier located behind the RSGW may receive thousand of reports from the ST. This is the "IGMP flooding" issue that causes a waste of bandwidth and a high CPU activity at the querier side.

- Latency Issue

Another essential issue mostly due to the IGMP behaviour is the latency when the last end-user leaves a multicast group. When the last member is leaving, IGMP implemented as defined in the RFC will take much time to detect that there is no member on the group anymore and also stop transmitting the multicast stream on the air interface. The time required is defined by the IGMP protocol and will lead to an important waste of bandwidth. IGMP max response time is set according to the estimated number of members in the group, but is should be tuned to ensure that few reports are received if there are still members and to reduce latency to leave if there is only one member remaining.

- IGMPv2 adaptation parameters

To solve previous issues, IGMP adaptation over satellite has been standardized under TS 102 293 [E] (V1.1.1). The adaptation enables to largely reduce the total volume of the IGMP traffic over BSM systems.

E.2 RSM-B IP multicast solution

RSM-B represents a BSM regenerative system based on a DVB-RCS / DVB-S on-board processing (OBP) satellite. RSM-B system is completed by a fleet of standard RCSTs, a MS (Management Station) acting as Network Control Centre, and several RSGWs (Regenerative Satellite GateWays) for interconnection with terrestrial networks. RSM-B supports star (RCST to RSGW) and mesh (RCST to RCST) connectivity in just one satellite hop with unidirectional or bi-directional, point-to-point and point to multipoint connections. The connection control, QoS assurance, addressing resolution and support of multicast require the usage of a signalling protocol between the satellite terminals and the NCC. The Connection Control Protocol (C2P) aims assessing and enhancing the control plane of RSM-B DVB-RCS system.

RSM-B ground segment IP multicast solution is based on IGMPv2 proxying at terminals and an IGMPv2 adapter on the regenerative gateway. As most mesh multicast services are expected to be multiconference type (controlled number of sources), and star multicast could be any source from terrestrial networks, only dynamic multicast group management is offered for star multicast. Upon contract service level, multicast transmission is enabled for one subscriber.

RSM-B system supports two types of IP multicast services based on two types of topologies:

- Star IP multicast.
- Mesh IP multicast.

In the Star IP multicast, multicast flows are dynamically forwarded from a RSGW to several RCSTs. Multicast sources are on the terrestrial network and forward their multicast flows towards the RSGW.

In the Mesh multicast, multicast flows are statically forwarded from a source RCST to several destinations RCSTs. Multicast sources are on terrestrial network and forward their multicast flows to a source RCST. Both services may be combined in the same RCST or RSGW.

When mapping to the different multicast scenarios proposed in the document, the two IP multicast services identified in RSM-B correspond to:

- Dynamic Star IP multicast to Star Pull Source via Hub.
- Static Mesh IP multicast to Mesh Push Source via ST.

E.2.1 Star IP Multicast

Star IP Multicast implies distribution of worldwide available IP multicast data flows if the RSGW has Internet Access and distribution of MSP, ISP or Corporate dedicated IP multicast data flows. The star dynamic multicast service allows any RCST that has access to a RSGW, to join any multicast source from terrestrial networks. Star multicast services are one to many only, such as NVoD, Broadcast TV/Radio and file transfer. These services are provided with a single hop satellite delay.

For star dynamic multicast (Figure E.2.1), multicast flows are dynamically forwarded from a RSGW to several RCSTs. The RSGW will forward into the satellite network the terrestrial network multicast sources under RCSTs request. The multicast forwarding is dynamic: the IGMPv2 protocol is running between the default RSGW and RCSTs to discover multicast group membership. A RSGW forwards only a multicast flow on the uplink assuming that at least one RCST has joined this flow, and the OBP replicates this flow to downlink TDMs covered by the RSGW. A RSGW does not forward multicast flow on the uplink if no RCST has joined it. The RSGW is in charge of managing the memberships so that to send each multicast flow only once.

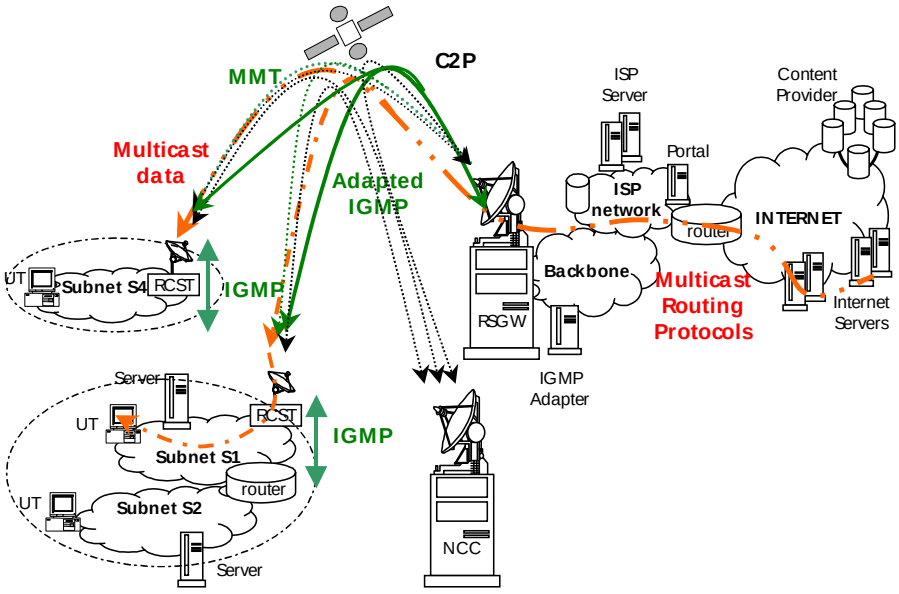

Figure E.2.1: RSM-B Star IP dynamic multicast

The network entities involved in the Star IP multicast network topology are hereafter overviewed:

- User Terminal (UT).

A User Terminal is on the IP subnet connected to the RCST through the User Interface. This UT has an IGMP v2 host function to subscribe/de-subscribe to a multicast group.

- Router on the IP subnet of a RCST.

There can be a router on the path between the RCST and User Terminal to connect several LANs to the RCST. This router with no PIM, support and behaves as an IGMP v2 Proxy based on the Internet Draft Draft-ietf-magma-igmp-proxy-06.txt .The upstream interface is the interface towards the RCST and the downstream interface is the interface towards User Terminals.

- RCST.

An IGMP v2 Proxy based on the Internet Draft Draft-ietf-magma-igmp-proxy-06.txt is implemented in a RCST.

The RCST acts as an IGMP Router and Querier on its User interface and it acts as an IGMP Host on the satellite air interface. The IGMP proxy avoids the RSGW receiving request message for an IP multicast flow from an UT, on the LAN of a RCST, which has already been requested by another UT on the same LAN. The IGMP Proxy knows that the IP multicast flow is already transmitted on the TDM and provides it to the new UT without asking the RSGW. This enables to reduce the number of IGMP messages sent on the air interface.

The IGMP proxy forwards multicast data flow received from the satellite air interface to its User Interface according to its group membership table.

- RSGW.

The RSGW is composed of:

- a GW_RCST;
- an IGMP adapter;
- a Multicast Edge Router.

The IGMP Adapter is an IGMP Proxy optimized to satellite environment following specification TS 102 293 [E]. The IGMP adapter is based on IGMP v2 proxy and performs specific functions to improve the signalling load induced by the IGMP v2 protocol, and to reduce the latency to stop multicast traffic for the satellite networks. It has an interface towards GW-RCSTs and an interface towards a Multicast Edge Router. The IGMP host function is standard as defined in RFC 2236 [T]. The IGMP Querier function is optimized to the satellite environment having variable values of timers, adapted to a satellite network.

The Multicast Edge Router is a multicast router with an IGMPv2 Router/Querier function on its interface towards the IGMP adapter and a multicast routing function on its other interface. The multicast routing function of the Multicast Edge Router is commonly based on PIM-SM. It is in charge of joining or pruning to the group-spanning tree.

All the IGMP messages are transparently forwarded no decrement of TTL (Time To Live) field value of IP packet, inside the RSGW till they reach the Multicast Edge Router. Therefore the GW_RCST transparently forwards all IGMP messages without decreasing the TTL field value.

- Backbone network, ISP or MSP network, Internet, Corporate Network.

These networks have Multicast Core Routers with multicast routing protocol such as PIM-SM, MBGP, and MSDP.

E.2.1.1 RCST Star IP Multicast functions

The IGMP Proxy included into the RCST has two interfaces. These interfaces must be configured as defined below:

- An IGMPv2 Host function on the satellite air interface.
- An IGMPv2 Querier on the User Interface.

Figure E.2.2 shows the RCST IGMP Proxying functionality:

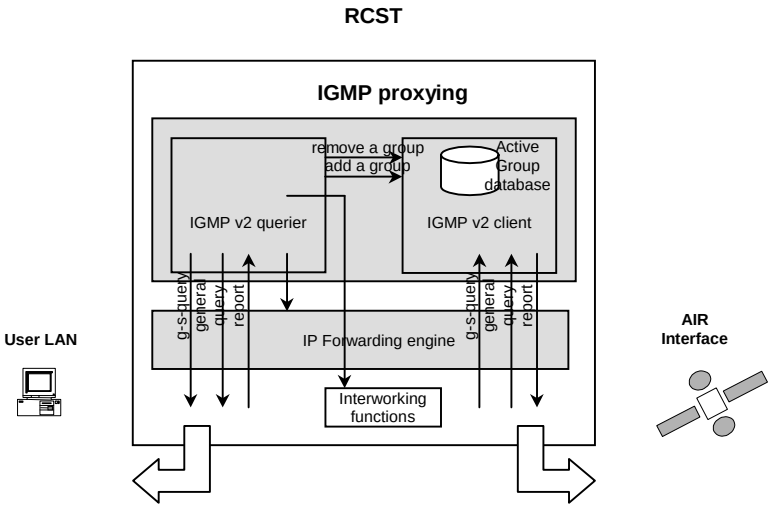

Figure E.2.2: IGMP proxying overview

The IGMPv2 reports sent by the RCST towards the RSGW have as IP source address the IP address of the RCST in its Air Interface.

The IGMP v2 proxy shall be enabled and configured to receive IGMP v2 Query messages from the multicast edge router (MER) of the RSGW. It shall process IGMPv2 messages using the multicast address ALL-SYSTEM (224.0.0.1) without sending an IGMP report for this group address towards the RSGW. This IGMP v2 proxy has a membership database consisting of the merge of all subscriptions on its downstream interface.

The IGMP Querier function on the User Interface is configured to be elected as Querier. This entity processes IGMP messages received from the User Interface.

The RCST star IP multicast functions follow the IGMP-specific BSM functional model defined in TS 102 293 [E] (V1.1.1).

E.2.1.1.1 RCST IGMPv2 Host

The IGMP v2 host shall process as specified in clause 6 of RFC 2236 [T]. It shall process one state diagram per multicast address active in the Membership database.

The IGMPv2 Host must not respond for a group with an address included in 224.0.0.0/24. These are local sub-network multicast addresses that do not need to be controlled by IGMP.

The "Join group" event appears on reception of an "add a group" from the IGMPv2 Querier. In addition to the specified actions on the "Join group" event, the IGMPv2 Host shall:

- Send an IGMP report to the RSGW. The unsolicited report must not be sent by the IGMPv2 Host if it receives a resent report from the IGMP Adapter before the unsolicited report timer expires.
- Add the group to the membership database.

The "Leave group" event appears on reception of a "remove a group" from the IGMPv2 Querier. In addition to the specified actions on the "Leave group" event, the IGMPv2 Host shall remove the group from the Membership database. The IGMPv2 Host shall send an IGMP Leave message if its flag is set. If the flag is not set, the IGMPv2 Host must not forward an IGMP leave message (silent leave).

When a General Query is received, the IGMPv2 Host shall act if it is a member of all groups present in the Membership database.

The unsolicited Report Interval of IGMPv2 Host shall be configurable.

Upon reception of a "leave" report from an RCST IGMP proxy, the RSGW multicast edge router will send a "Group Specific Query" for the corresponding multicast group, that will be replied with a membership report from one of the RCST IGMP proxys (if any) listening to the session.

E.2.1.1.2 RCST IGMPv2 Querier

The IGMPv2 Querier shall process clause 7 of RFC 2236 [T] with the following extra actions:

On the "notify routing +" action, the IGMPv2 Querier shall:

- Change the group membership table used to forward multicast packets.
- Send a "add a group" to the IGMPv2 Host.

On the "notify routing -" action, the IGMPv2 Querier shall:

- Change the Group Membership table used to forward multicast packets.
- Send a "remove a group" to the IGMPv2 Host.

Some parameters of the IGMPv2 Querier shall be configurable:

- Robustness variable.
- Query Interval.
- Query Response Interval.

To receive IGMP messages and Multicast Traffic sent by the RSGW, the RCST needs to filter the multicast PID associated to these data. The multicast PID is indicated into the Multicast Mapping Table (MMT) forwarded by the NCC. The RCST keeps in memory a copy of the last received MMT. The RCST learns the PID use to decode the MMT from the Network Layer Information Descriptor received in the TIM_u during the logon process.

The RCST shall first learn the PID used to forward IGMP messages (IP address 224.0.0.1). After upon creation of a group membership (first IGMP Report message on its User Interface) the IGMPv2 Proxy shall send an internal request to check the corresponding multicast PID in the MMT and then updates the list of PIDs to listen to. Upon suppression of a group membership, the IGMP Proxy shall send an internal request to remove the corresponding PID from the list of PIDs to listen to.

To check the multicast PID corresponding to an IP multicast address, the RCST shall parse the MMT from the beginning to end and shall select the first PID corresponding to the IP multicast address.

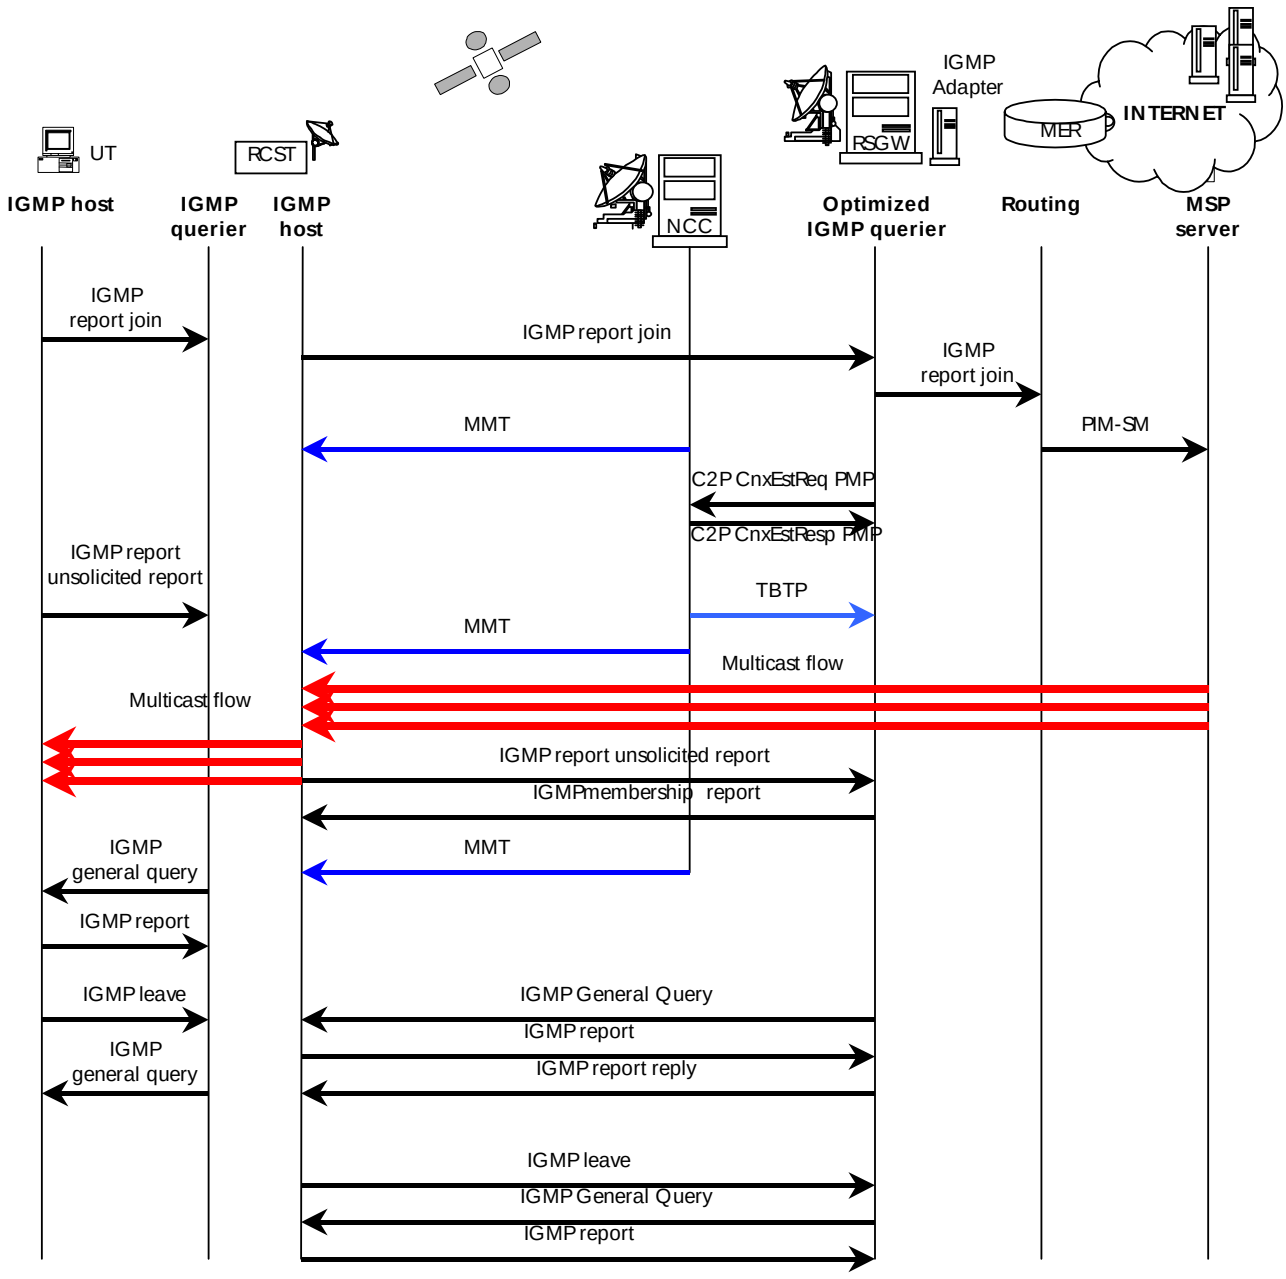

Figure E.2.3: Star IP Multicast scenario example

E.2.1.2 Connections for Star IP Multicast

The NCC is in charge of assigning a multicast PID to the RSGW and of periodically sending the Multicast Map Table of a satellite network. A typical refresh period for the MMT is 10 s. This MMT is updated according to RSGW request (Connection Establishment Request for a point-to-multipoint connection).

IGMP v2 protocol needs a two-way communication:

- The RSGW sends IGMP v2 messages to all covered TDMs.
- The RCSTs send IGMP v2 messages to the RSGW.

Multicast flows need a unidirectional point-to-multipoint connection from the RSGW to all covered TDMs. Then the following connections are set-up:

- A point-to-multipoint connection from the RSGW to all covered TDMs to forward IGMP v2 messages and all multicast flows.

- A point-to-point connection from each RCST to the RSGW to forward IGMP v2 messages.

E.2.2 Mesh IP Multicast

The mesh static multicast service provides forwarding of multicast data between RCSTs having LAN interconnection over the RSM-B system. Mesh multicast services (from end-users to end-users) such as multi-video-conferencing and shared workspace are provided with as in the star mode, in single hop satellite delay (thanks to RSM-B OBP).

In meshed static multicast (Figure E.2.4), multicast flows are statically forwarded from a source RCST to several destination RCSTs. The OBP replicates this flow to all downlink TDMs where there could be RCSTs belonging to the same satellite network. Multicast sources are hosted by the RCST which is in charge of forwarding the multicast flows to other RCSTs through the satellite interface. It is the same network architecture than in Star IP Multicast, except that only RCSTs are involved.

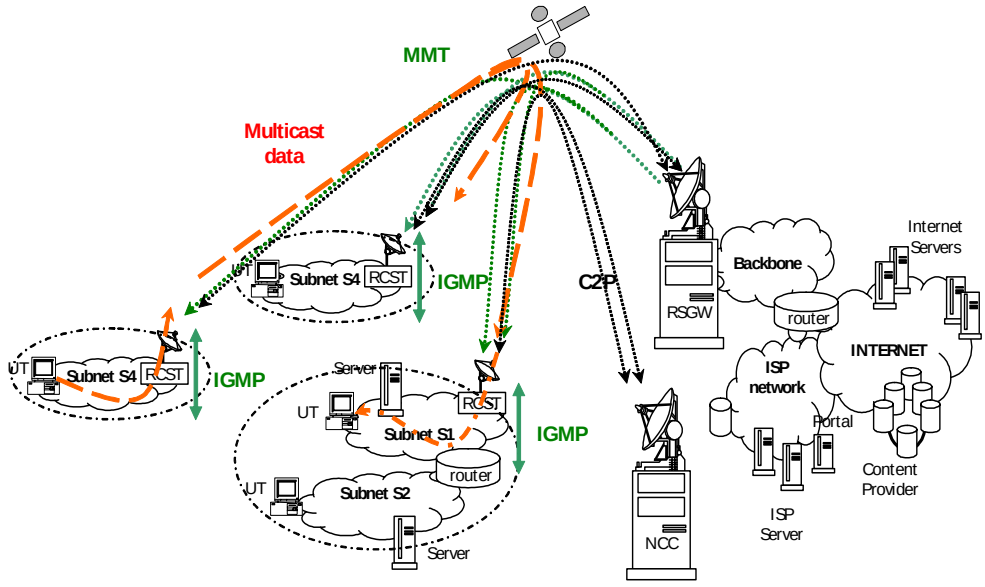

Figure E.2.4: Mesh IP multicast network topology

The Network entities involved in RSM-B mesh multicast solution are hereafter reviewed:

- User Terminal (UT).

A User Terminal is on the IP subnet connected to the RCST through the User Interface. This UT has an IGMP v2 host function to subscribe/un-subscribe to a multicast group.

- Router on the IP subnet of a RCST.

There can be a router on the path between the RCST and User Terminal to allow several LANs connected to the same RCST. This router with no PIM support, behaves as an IGMP v2 Proxy based on the Internet Draft Draft-ietf-magma-igmp-proxy-04.txt. The upstream interface is the interface towards the RCST and the downstream interface is the interface towards User Terminals.

Moreover this router is configured to forward some multicast flows to the RCST. The list of authorized multicast flows is the same as the one configured into the RCSTs.

- RCST.

The RCST is an IGMP v2 Querier to process subscription of UTs on the User Interface. On the air satellite interface a RCST has no IGMP function. It forwards multicast data flow received from the air satellite interface to its User Interface according to its group membership table when requested.

In addition a RCST has the list of IP multicast group addresses that authorized to be forwarded to the air satellite interface. This list is defined per RCST and is configured by management. An RCST forwards these authorized IP multicast flows from the User Interface to the air satellite interface over a point-to-multipoint connection.

- RSGW.

The RSGW will allow reception of the mesh multicast flows and forwarding them to the multicast bone when requested. To accomplish this task, the multicast edge router will support multicast border protocols such as PIM-SM, MBGP and MSDP.

- Multicast addresses.

Each satellite network managed by a service provider has a pool of IP multicast addresses assigned as recommended in the RFC 2365 [W]. Each RCST has pool of IP multicast addresses authorized to be forwarded.

E.2.2.1 RCST Mesh IP Multicast functions

Mesh IP Multicast functions require an IGMPv2 querier implemented in the RCST. This RCST querier implementation follows RFC 2236 [T].

An overview of IGMPv2 querier function is given in figure E.2.5:

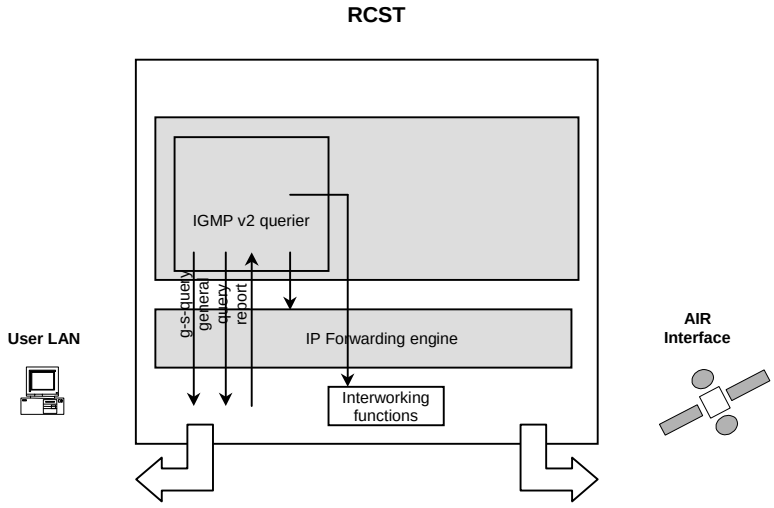

Figure E.2.5: IGMPv2 querier overview

To receive multicast traffic sent by a RCST, the RCST needs to filter multicast PIDs associated to the traffic. These PIDs are indicated into the Multicast Map Table (MMT) forwarded by the NCC.

The NCC periodically sends a MMT. The MMT is described in the TR 101 790 [U] clause I.6 and gives the multicast PID associated to a multicast IP address. Then the RCST processes the MMT and keeps in memory a copy of the last received MMT.

Upon creation of a group membership (first IGMP Report message on its User Interface) the IGMP Querier sends an internal request to look up the corresponding multicast PID in the MMT and then updates the list of PIDs to listen to. Upon suppression of a group membership, the IGMP Querier sends an internal request to retrieve the corresponding PID from the list of PIDs to listen to.

Each RCST may transmit and receive IP multicast using the same IP multicast address but using different PIDs. MMT must be fully scanned to extract all the PIDs linked to the same IP multicast address.

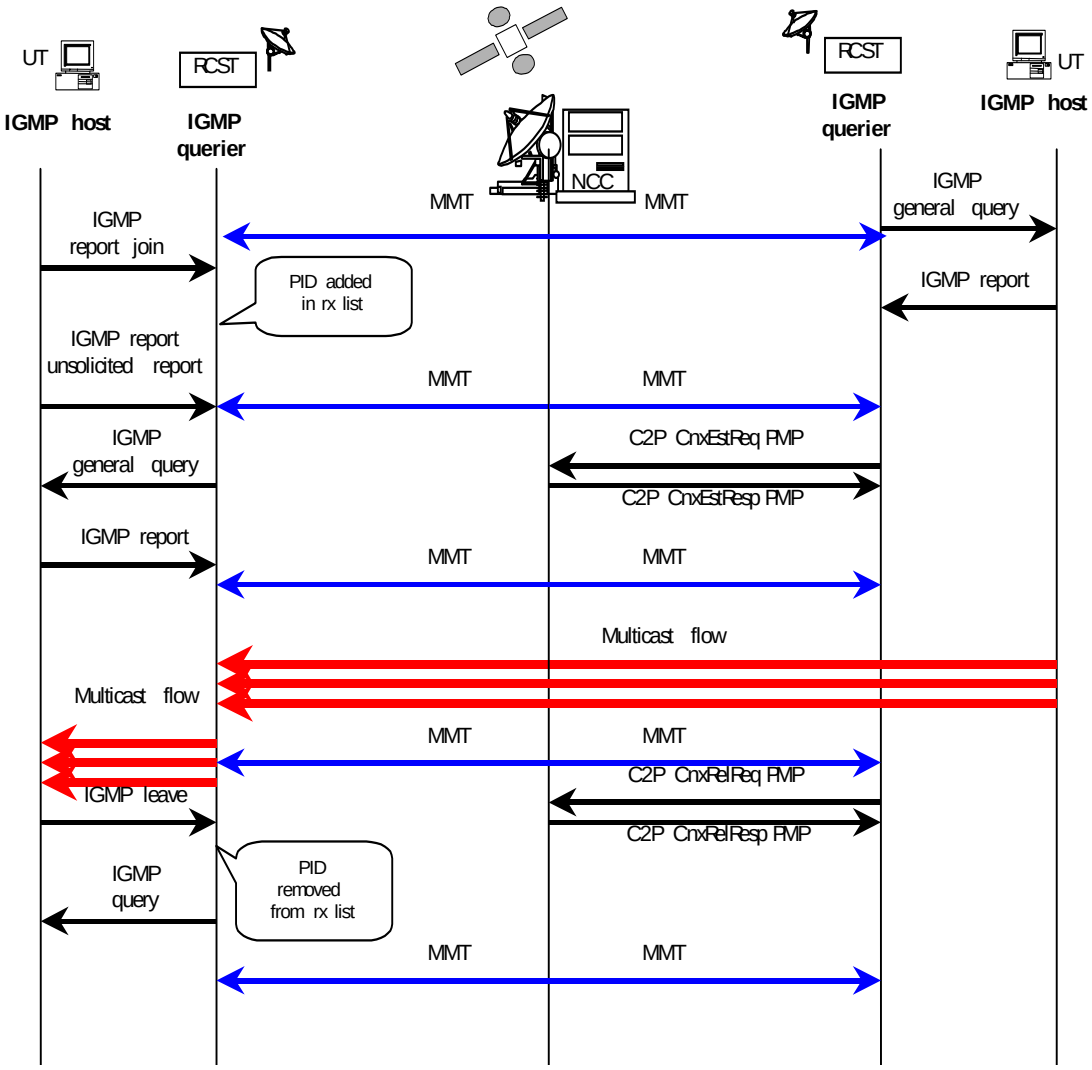

Figure E.2.6: Mesh IP Multicast scenario example

Annex F (informative): Bibliography

[A] ETSI TR 101 984: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; Services and Architectures".

[B] ETSI TR 101 985: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; IP over Satellite".

[C] ETSI TR 102 156: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; IP interworking over satellite; Multicasting".

[D] ETSI TS 102 292: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM) services and architectures; Functional architecture for IP interworking with BSM networks".

[E] ETSI TS 102 293: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; IP Interworking over satellite; Multicast Group Management; IGMP adaptation".

[F] ETSI TS 102 460: "Services and Architectures: Address Management at the SI-SAP".

[G] ETSI TS 102 463: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM); Interworking with IntServ QoS".

[H] ETSI TS 102 464: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM); Interworking with DiffServ Qos".

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

[J] IETF RFC 3754: "IP Multicast in Differentiated Services (DS) Networks" .

[K] IETF RFC 3171: "IANA Guidelines for IPv4 Multicast Address Assignments" .

[L] IETF RFC 4541: "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches".

[M] IETF RFC 4605: "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")".

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

[O] IETF RFC 4608: "Source-Specific Protocol Independent Multicast in 232/8".

[P] IANA, "Address Family Numbers", http://www.iana.org/assignments/address-family-numbers.

[Q] draft-ietf-ipdvb-ar-xx.txt, "Address Resolution Mechanisms for IP Datagrams over MPEG-2 Networks", G. Fairhurst, M. Montpetit, IETF ipdvb WG, 2007 (scheduled for RFC Publication as an informational document, in early 2007)

[R] IETF RFC 4259: "A Framework for Transmission of IP Datagrams over MPEG-2 Networks".

[S] IETF RFC 2464: " Transmission of IPv6 Packets over Ethernet Networks ".

[T] IETF RFC 2236: "Internet Group Management Protocol, Version 2".

[U] ETSI TR 101 790: "Digital Video Broadcasting (DVB); Interaction channel for Satellite Distribution Systems; Guidelines for the use of EN 301 790".

[V] IETF RFC 3077: "A Link-Layer Tunneling Mechanism for Unidirectional Links ".

[W] IETF RFC 2365: "Administratively Scoped IP Multicast".

[X] IETF RFC 2375: "IPv6 Multicast Address Assignments".

[Y] IEEE 802: "Local and Metropolitan Area Network Standards".

[Z] IETF RFC 1701: "Generic Routing Encapsulation (GRE)".

[AA] IETF RFC 1702: "Generic Routing Encapsulation over IPv4 networks".

History

Document history		
V1.1.1	January 2007	Publication