ETSI TR 102 156 V1.1.1 (2003-04)

Technical Report

Satellite Earth Stations and Systems (SES);
Broadband Satellite Multimedia;
IP interworking over satellite;
Multicasting

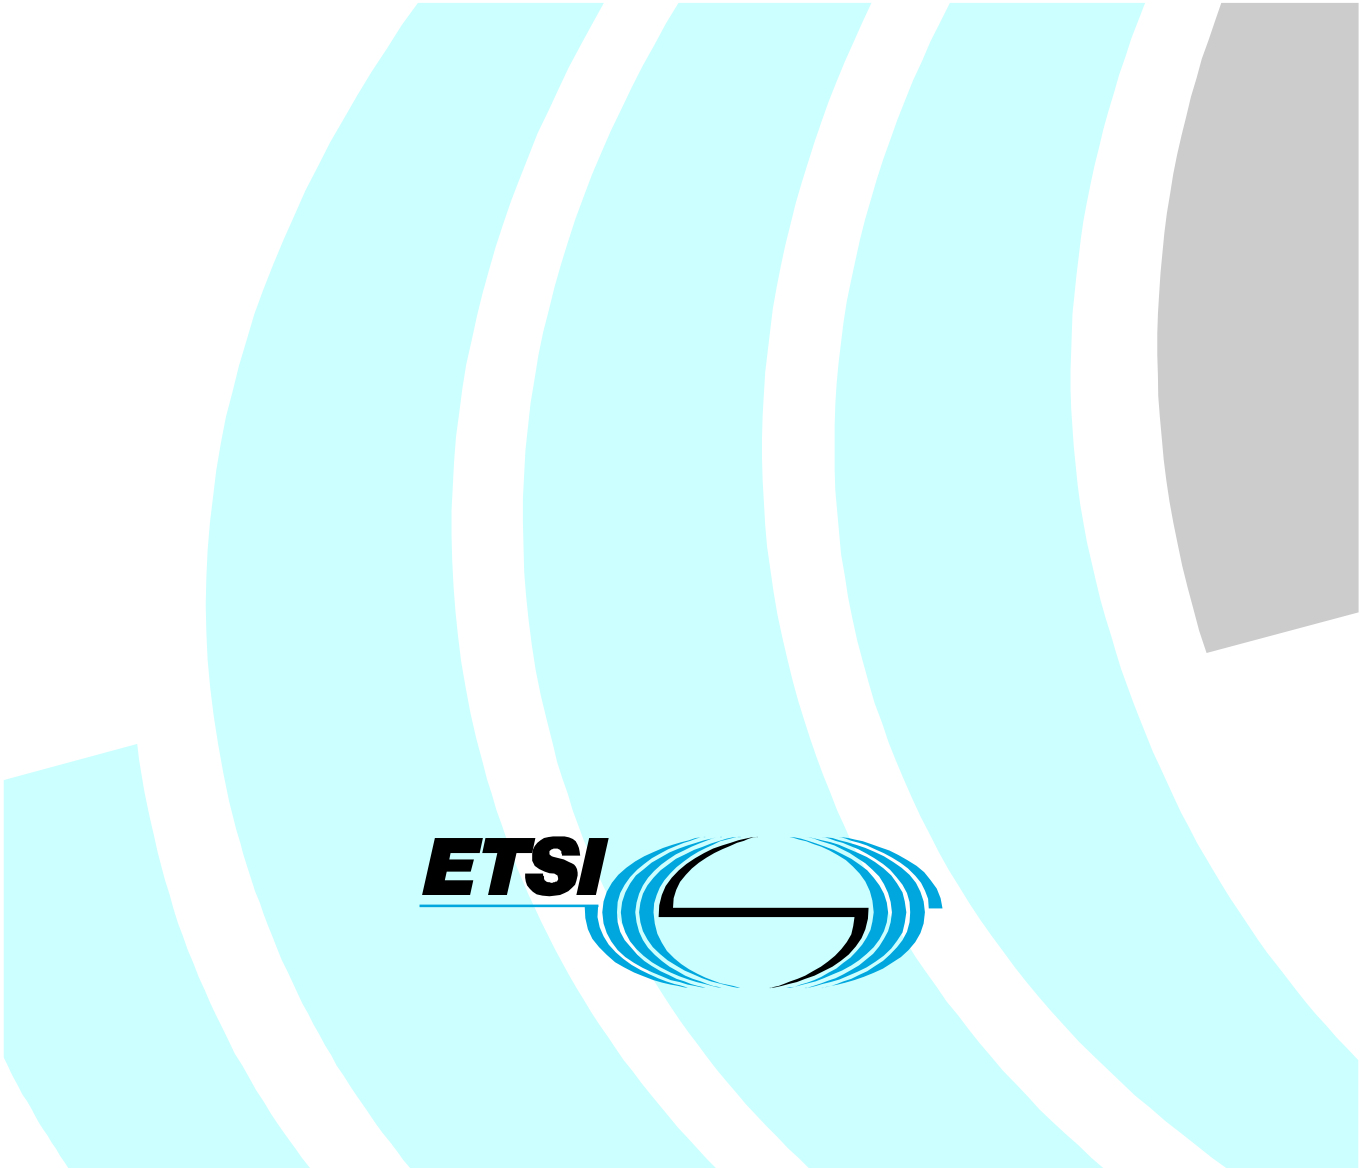

Reference
DTR/SES-00075

Keywords
interworking, broadband, IP, multimedia, satellite

ETSI

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

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

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

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).

All published ETSI deliverables shall include information which directs the reader to the above source of information.

Foreword

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

The present document has been generated by ETSI Specialist Task Force STF 214 "Broadband and Satellite integration" [5].

Introduction

This work is based on the two ETSI Technical Reports, TR 101 984 [4] and TR 101 985 [3].

The present document identifies and discusses issues related to a standardization of Broadband Satellite Multimedia (BSM) multicast, and suggests technical specifications to be produced by ETSI.

The present document focuses on proposed standards for delivering IP Multicast over BSM networks, since IP is expected to be the most common multicast technology. However, we recognize that satellite multicast concepts do need not be based on IP, but non-IP multicast presently lack a portfolio of protocols and applications compared to IP.

It is assumed that the reader is somewhat familiar with IP, multicast concepts and BSM satellite systems.

1 Scope

The focus of the present document is satellite-based multicasting, including IP multicasting.

The scope of the present document is to:

- identify relevant multicast issues, use cases service architectures for satellite multicasting;
- identify satellite specific issues and technical requirements for satellite multicasting;
- identify relevant standardization work in other standards bodies;
- and conclude what actions ETSI should be taking with respect to preparing Technical Specifications.

2 References

For the purposes of this Technical Report (TR) the following references apply.

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

[1] ETSI TR 101 374-1: "Satellite Earth Stations and Systems (SES); Broadband satellite multimedia; Part 1: Survey on standardization objectives", 1998.

[2] ETSI TR 101 374-2: "Satellite Earth Stations and Systems (SES); Broadband satellite multimedia; Part 2: Scenario for standardization", 2000.

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

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

[5] STF 214 Terms of Reference; Terms of Reference for Specialist Task Force STF 214 (SES) on Broadband Satellite Internet Integration; 2001.

[6] IETF RFC 2887 (2000): "The Reliable Multicast Design Space for Bulk Data Transfer".

[7] IETF RFC 1112 (1989): "Host Extensions for IP Multicasting".

[8] IETF RFC 2236 (1997): "Internet Group Management Protocol, Version 2".

[9] IETF RFC 2326 (1998): "Real Time Streaming Protocol (RTSP)".

[10] IETF RFC 2117 (1997): "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification".

[11] IETF RFC 1889 (1996): "RTP: A Transport Protocol for Real-Time Applications".

[12] IETF RFC 2365 (1998): "Administratively Scoped IP Multicast".

[13] IETF RFC 1884 (1995): "IP Version 6 Addressing Architecture".

[14] IETF RFC 1458 (1993): "Requirements for Multicast Protocols".

[15] IETF RFC 2627 (1999): "Key Management for Multicast: Issues and Architectures".

[16] IETF RFC 3307 (2002): "Allocation Guidelines for IPv6 Multicast Addresses".

[17] IETF RFC 2715 (1999): "Interoperability Rules for Multicast Routing Protocols".

[18] IETF RFC 3353 (2002): "Overview of IP Multicast in a Multi-Protocol Label Switching (MPLS) Environment".

[19] M. Koyabe and G. Fairhurst: "Performance of reliable multicast protocols via satellite at EHF persistent fades", Report, Electronics Research Group University of Aberdeen King's College, Old Aberdeen, Aberdeen, AB24 3UE Scotland.

[20] Mohammad Banikazemi: "IP Multicasting: Concepts, Algorithms, and Protocols, IGMP, RPM, CBT, DVMRP, MOSPF, PIM, MBONE".

[21] M.Koyabe and G. Fairhurst: "Reliable Multicast Via Satellite: A Comparison Survey and Taxonomy", International Journal for Satellite Communications (IJSC), Vol:24(1), 21-26, 2001.

[22] IETF RFC 2362 (1998): "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification".

[23] IETF RFC 1075 (1988): "Distance Vector Multicast Routing Protocol".

[24] IETF RFC 3376 (2002): "Internet Group Management Protocol, Version 3".

[25] IETF RFC 2093 (1997): "Group Key Management Protocol (GKMP) Specification".

[26] IETF RFC 2094 (1997): "Group Key Management Protocol (GKMP) Architecture".

[27] IETF RFC 1949 (1996): "Scalable Multicast Key Distribution".

[28] Draft-ietf-rmt-info-fec-03 (2003): "The Use of Forward Error Correction in Reliable Multicast".

[29] IETF RFC 3077 (2001): "A Link-Layer Tunneling Mechanism for Unidirectional Links".

[30] IETF RFC 1119 (1989): "Network Time Protocol (version 2) specification and implementation".

[31] ETSI TR 102 157: "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; IP Interworking over satellite; Performance, Availability and Quality of Service".

[32] ITU-T Recommendation I.571: "Connection of VSAT based private networks to the public ISDN".

[33] ITU-T Recommendation I.572: "VSAT interconnection with the PSTN".

[34] ETSI EN 301 790: "Digital Video Broadcasting (DVB); Interaction channel for satellite distribution systems".

[35] Draft-ietf-msdp-spec-14 (2002): "Multicast Source Discovery Protocol (MSDP)".

[36] Draft-ietf-ssm-overview-04 (2002): "An Overview of Source-Specific Multicast (SSM)".

[37] IETF RFC 2933 (2000): "Internet Group Management Protocol MIB".

[38] Draft-ietf-magma-igmp-proxy-02 (2002): "IGMP/MLD-based Multicast Forwarding ("IGMP/MLD Proxying")".

3 Definitions and abbreviations

3.1 Definitions

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

gateway: network point that acts as an entrance to another network; often a connection point between a satellite network and the core network

NOTE: Traffic exits to the core network via a gateway when it originates in the satellite network.

satellite access function/terminal: end point for a satellite connection that does not connect to the core network

NOTE: In a mesh network terminals and gateways can in principle be the same hardware with equal air interfaces, but will normally be distinguishable by their terrestrial connections and their size/capacity.

Multicast Entry Point (MEP): point of entry into the BSM network for an external multicast source

NOTE: Does not need to be a fixed entry point. Temporary definition for TR 102 156.

Network Entry Point (NEP): point of entry for external communication into the BSM network

Network Control Center (NCC): satellite network node that controls the network resources

NOTE: May hold subscriber data and info like where users or terminals are at any time.

host/multicast host: IETF terminology for any computer or network device that has full two-way access to other computers on the Internet

NOTE: In the context of multicast it is a computer that sends or receives multicast data.

multicast: communication capability, which denotes unidirectional distribution from a single source to a number of specified destination points

NOTE: Transmission from one terminal or gateway to one or more specified destinations.

unicast/point-to-point: transmission from one terminal or gateway to another

anycast: communication between a single sender and the nearest of several receivers in a group (in IPv6)

broadcast: Communication capability which denotes unidirectional distribution to an unspecified number of access points connected to the network. Transmission to all network destinations within an IP subnetwork.

intra-domain: communications within a internet subdomain, e.g. within the BSM network

inter-domain: communications between to or more different domains, e.g. the BSM network and other terrestrial networks

session: In this usage, the session is regarded as the logical setup. A session could be terminated and the connection maintained for a new session later.

connection: the physical setup so two or more hosts can communicate

session: the logical setup so two or more hosts can communicate

streaming: streaming media includes e.g. streaming video with sound

NOTE: The media is sent in a continuous stream and is played as it arrives.

On-Board Processing (OBP): satellite capability to receive and process data, and an ability to switch or route data, possibly replicate and communicate directly from terminal to terminal

bent pipe satellite: satellites act as a frequency translator, and transmit exactly the same data they receive, just shifted to another frequency

open groups: publicly available multicast groups, that generally can be subscribed to both within and outside the BSM network

closed groups: multicast groups that are available only within the BSM network

channel: means of unidirectional transmission of signals between two points

NOTE: Several channels may share a common transport mechanism.

management plane: the management plane provides two types of functions, namely Layer Management and plane management functions:

- **user plane:** the user plane has a layered structure and provides for user information flow transfer and associated controls (e.g. flow control, recovery from errors, etc);
- **control plane:** the control plane 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.

3.2 Abbreviations

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

AAL	ATM Adaptation Layer
ACK	ACKnowledgement
AFDP	Adaptive File Distribution Protocol
API	Application Programming Interface
ARP	Address Resolution Protocol
ARQ	Automatic Repeat reQuest
AS	Autonomous System
ASM	Any Source Multicast
AVT	Audio/Video Transport
BGMP	Border Gateway Multicast Protocol
BGP	Border Gateway Protocol
BR	Border Router
BSM	Broadband Satellite Multimedia
BSM-M	BSM Multicast
BSM-MP	BSM Multicast Proxy
BSR	BootStrap Router
CA	Conditional Access
CAT	Conditional Access Table
CBT	Core Based Trees
CDMA	Code Division Multiple Access
CP	Contact Point
CP	Customer Premises
CW	control word
DHCP	Dynamic Host Configuration Protocol
Diffserv	Differentiated services
DM	Dense Mode
DNS	Domain Name System
DR	Designated Router
DVB	Digital Video Broadcasting
DVB-RCS	DVB Return Channel System
DVMRP	Distance Vector Multicast Routing Protocol
ECM	Entitlement Control Message
ECTP	Enhanced Communications Transport Protocol
EMM	Entitlement Management Messages
EU	European Union
FCAPS	Fault, Configuration, Authentication, Performance, Security
FDDI	Fiber Distributed Data Interface
FEC	Forward Error Correction
FTP	File Transfer Protocol
GEO	Geostationary Earth Orbit
GKMP	Group Key Management Protocol
GMP	Group Management Protocol
GSM	Global System for Mobile telephony
HLR	Home Location Register
HTTP	HyperText Transfer Protocol
IANA	Internet Assigned Numbers Authority
ICEBERGS	IP ConferEncing with Broadband multimedia ovERGeostationary Satellites
ICMP	Internet Control Message Protocol
IDMR	Inter-Domain Multicast Remnants
IEC	International Electrotechnical Commission
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IP	Internet Protocol
IPv4	Version 4 (current) of the Internet Protocol
IPv6	Version 6 (next generation) of the Internet Protocol
ISL	Inter-Satellite Link

ISP	Internet Service Provider
ITU	International Telecommunications Union
LAN	Local Area Network
LBRM	Log-Based Receiver Reliable Multicast
LGMP	Local Group based Multicast Protocol
LLTM	Link-Layer Tunneling Mechanism
LSMA	Large Scale Multicast Applications
MAC	Media Access Control
MAC	Message Authentication Code
MALLOC	Multicast address ALLOCation
MASC	Multicast Address-Set Claim
MBGP	Multiprotocol Border Gateway Protocol
MBone	Multicast Backbone
MBoneD	MBone Deployment
MCS	Master Control Station
MCU	Multipoint Control Unit
MDP	Multicast Dissemination Protocol
MEP	Multicast Entry Point
MF/TDMA	Multi Frequency Time Division Multiple Access
MFTP	Multicast File Transfer Protocol
MGID	Multicast Group ID
MIB	Management Information Base
MMT	Multicast Mapping Table
MOSPF	Multicast Open Shortest Path First
MP	Multicast Proxy
MPE	Multi Protocol Encapsulation
MPLS	MultiProtocol Label Switching
MR	Multicast Router
MRIB	Multicast RPF Routing Information Base
MSC	Mobile Service Center
MSDP	Multicast Source Discovery Protocol
MTP	Multicast Transport Protocol
NACK	Negative ACKnowledgement
NCC	Network Control Centre
NEP	Network Entry Point
NFS	Network File System
NGN	Next Generation Networks
NOC	Network Operation Centre
NSP	Network Service Provider
NTP	Network Time Protocol
OAM	Operations And Maintenance
OBP	OnBoard Processing
OSI	Open Systems Interconnection
OTERS	On-Tree Efficient Recovery using Subcasting
PGM	Pragmatic General Multicast
PID	DVB Packet Identifier
PIM	Protocol Independent Multicast
PMT	Program Map Table
PTV	Pay TV
QoS	Quality of Service
RAMP	Reliable Adaptive Multicast Protocol
RCMBS	Reliable Concurrent Multicast from Bursty Sources
RFC	Request For Comments
RM	Reliable Multicast
RMDP	Reliable Multicast data Distribution Protocol
RMP	Reliable Multicast Protocol
RMT	Reliable Multicast Transport
RMTP	Reliable Multicast Transport Protocol
RP	Rendezvous Point
RPF	Reverse Path Forwarding
RRM	Radio Resource Management
RRMP	Restricted Reliable Multicast Protocol

RSVP	ReSerVation Protocol
RTCP	Real Time Control Protocol
RTP	Real-time Transport Protocol
RTSP	Real Time Streaming Protocol
SA	Source-Active
SACK	Selective ACKnowledgement
SAP	Satellite Access Point
SAS	Subscriber Authorization System
SDP	Session Description Protocol
SFM	Source Filtered Multicast
SG	Study Group
SIM	Subscriber Identification Module
SIP	Session Initiation Protocol
SIP	Service Independent Protocol
SI-SAP	Satellite Independent-Service Access Point
SIT	Satellite Interactive Terminal
SKC	Session Key Changeover
SKE	Session Key Exchange
SM	Sparse Mode
SME	Security Message Exchange
SMS	Subscriber Management System
SOC	Satellite Operations Center
SOHO	Small Office Home Office
SRM	Scalable Reliable Multicast
SSM	Source Specific Multicast
ST	Satellite Terminal
STB	Set-Top Box
TCP	Transmission Control Protocol
TDM	Time Division Multiplexing
TDMA	Time Division Multiple Access
TIA	Telecommunications Industry Association (US)
TMTP	Tree-based Multicast Transport Protocol
TR	Technical Report
TRAM	Tree-based Reliable Multicast
TRM	Transport protocol for Reliable Multicast
TTL	Time To Live
UA	User Agent
UDL	UniDirectional Links
UDLR	UniDirectional Link Routing
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunications Service
UNI	User to Network Interface
UR	Unicast Router
VCC	Virtual Channel Connection
VOD	Video On Demand
VPC	Virtual Path Connection
VPN	Virtual Private Network
WAN	Wide Area Network
WG	Working Group
WLAN	Wireless LAN
WWW	World Wide Web

4 Multicasting background

The purpose of this clause is to present relevant, but non-satellite specific, multicast standards, procedures and terminology. Readers familiar with these concepts may skip the clause.

4.1 Introduction

Multicasting implies that addressed data is transmitted to several recipients with a single communications process. Multicast has a single stream of data for many users, in contrast to unicast technology, which employs a separate stream for each user, and broadcast, where all receivers receive the same non-addressed content. Multicast technology can reduce traffic on a network and on servers by eliminating redundant access to the same content.

Multicast technology can be used to send common data - such as streaming media as in a web concert or software updates, common information such as stock quotes or database or inventory updates - simultaneously from one-to-many or even many-to-many sources, as in server to client or application to application. Multicast is an ideal technology for enabling communication applications such as company meetings, distance learning, real-time training, and group conferencing.

Satellites have an architecture that lends itself to efficient broadcast and multicast applications. A satellite footprint can cover a very large geographical region. The cost of using satellite communications drops when a link can be shared with others. We can assume that for satellites the following apply:

- IP multicast content will globally cover a large number of interest groups. This is true also for users on satellite systems, as they could include users over a large area.
  - However, the receiver *density* on a given satellite system for a given source may be relatively low.
  - Open IP multicast groups will in general have the majority of hosts *outside* any specific satellite network.
- A BSM network may additionally offer restricted/private groups, i.e. for corporate multicasts and Virtual Private Networks (VPN).
  - Satellite service providers may also offer limited sources and groups just for the BSM subscribers, and such local sources may take particular advantage of BSM network capabilities as one in this case could have all receivers on the same subnetwork.

Satellite networks characteristics are influenced by delays corresponding to path delays, usually as for GEO satellites. Satellite transmission can also, as all wireless technologies, suffer from transmission errors due to various natural causes on the channel.

There will be different implementations of Broadband Satellite Multimedia systems. However, from a high level perspective, a BSM system will generally be composed of three segments as follows:

1) User segment;
   - comprising several types of terminals providing satellite access functions.
2) Space segment;
   - comprising one or several satellites, a Satellite Control Centre.
3) Operator segment;
   - comprising one or more Network Control Centers (NCC) and satellite gateways that interconnect to a terrestrial telecommunications infrastructure (e.g. to a terrestrial core network);
   - the logical multicast entry (and exit) point will for external sources be in this segment.

The source for the multicast can be either internal or external to the BSM network, as indicated in figure 1, and the content must be delivered to a set of addressed receivers, targeted at a set of different end-computers, or hosts, as they are often called in the internet language. A very general diagram showing two paths from external and internal sources to the hosts is indicated below. The figure shows that for an external source, it will enter the BSM network at some point, which then will be the Multicast Entry Point, MEP, and flow to the satellite via an uplink function before reaching a satellite terminal at the downlink, where it will be delivered to the particular hosts that have requested it. The figure also shows that a source can originate within the BSM network either in the operator segment or in the user segment from one of the hosts connected to the BSM network (lower left host acts as a source).

This is basic BSM multicast, and the challenge is to find the best way to handle this type of traffic in a BSM satellite network.

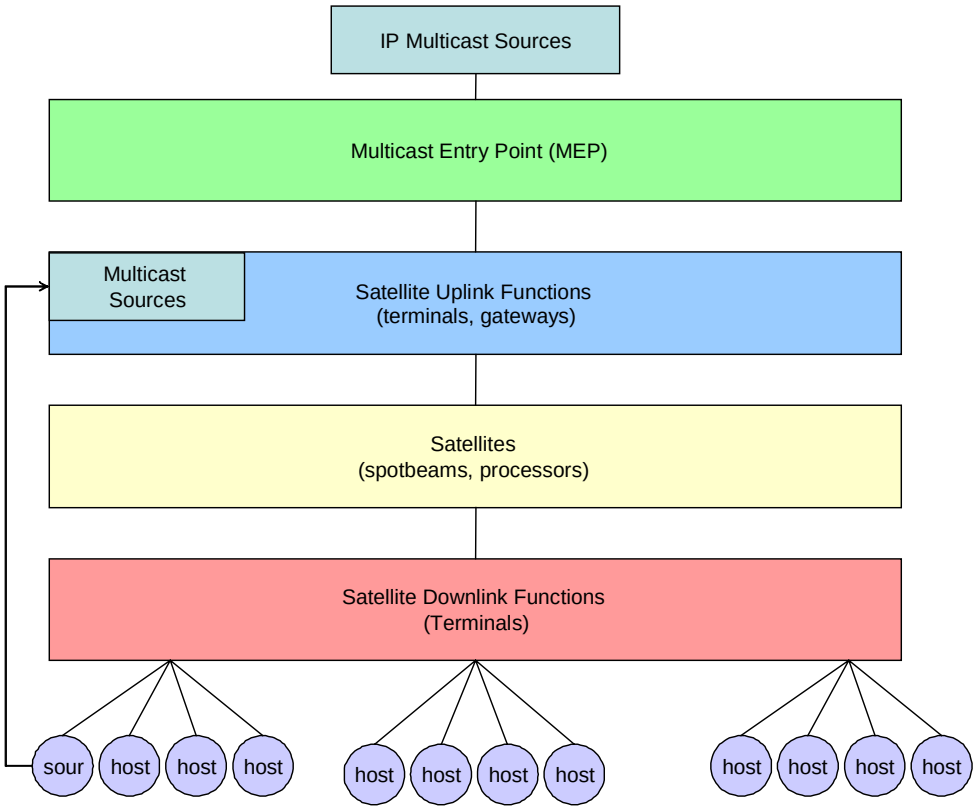

Figure 1: The satellite multicast challenge

A BSM multicast concept should offer both a well-defined inter-domain multicast interface to satellite terminals and gateways, and a well defined intra-domain multicast behavior. Sources can be either external Internet sources, internal form hosts connected to satellite terminals or closed BSM content added for instance in the gateways.

Some applications may require real-time content delivery, as in the case of interactive applications or time-sensitive applications, but most multicast networks today only support best-effort approaches. Real time applications may include interactive video-conferencing to many parties (teaching, learning), while non-real time may include video-on-demand or web-casting (e.g. a web-concert). There may however be time limitations to an application that is not real time. A video decoder may not be able to use data for an old picture frame if received too late relative to when the last data was received. Finally, multicasting data may also have requirements on time. For instance could it be required that messages are delivered to all receivers within a small time period, or within a given time interval, e.g. for stock quotes, or for some interactive gaming. It may also be a requirement to deliver data to users before the data has a certain age.

4.2 IP multicasting

Multicast is communication between (usually) a single sender and multiple receivers on a network. *Internet Protocol (IP) multicast* is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients and homes. Applications that take advantage of multicast include videoconferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news.

The first IP specification (IETF RFC 1112 [7]) on multicasting (Host Extensions for IP multicasting) was adopted by the Internet Engineering Task Force (IETF) in March of 1992 as the standard protocol for building multicast applications on the Internet. IP multicasting as specified in IETF RFC 1112 [7] is the transmission of an IP datagram to a set of zero or more hosts identified by a single IP destination address. A multicast datagram is delivered to all members of its destination host group with the same reliability as regular unicast IP datagrams.

There are three types of IPv4 addresses: unicast, broadcast, and multicast. Unicast addresses are used for transmitting a message to a single destination node. Broadcast addresses are used when a message is supposed to be transmitted to all nodes in a subnetwork. For delivering a message to a group of destination nodes which are not necessarily in the same subnetwork, multicast addresses are used. Whenever a multicast router receive a multicast packet it checks the group ID of the message and forwards the packet only if there is a member of that group in networks connected to it.

The membership of a host group is dynamic; and hosts may join and leave groups at any time without any restriction on the location or number of members in a host group. A host may be a member of more than one group at a time, and a host need not in principle be a member of a group to send datagrams to it.

A host group may be permanent or transient.

- A permanent group has a well-known, administratively assigned IP address. It is the address, not the membership of the group that is permanent. A permanent group may at any time have any number of members, even zero.

- Transient groups exist only as long as they have members Those IP multicast addresses that are not reserved for permanent groups are available for dynamic assignment to transient groups.

Forwarding of IP multicast datagrams is handled by "*multicast routers*". A host transmits an IP multicast datagram as a local network multicast which reaches all immediately-neighboring members of the destination host group. The multicast router(s) attached to the local network take responsibility for forwarding it towards other networks that have members of the destination group and so on until the destinations are reached.

There are three levels of conformance to IETF RFC 1112 [7]:

- Level 0: no support for IP multicasting.

- Level 1: support for sending but not receiving multicast IP datagrams.

- Level 2: full support for IP multicasting.

IP/Multicast is an extension to the standard Ipv4 network-level protocol that supports multicast traffic. Multicast is an option in IPv4, but a standard feature of IPv6.

Multicast enabled IP routers organize each multicast group into a spanning tree, and route multicast packets by making a copy of each packet for each output interface that includes at least one downstream member of the multicast group.

Multicasting is considerably more efficient when a subnetwork explicitly supports it (IETF pilc WG and [22]). For example, a router relaying a multicast packet onto an Ethernet subnet need send only one copy, no matter how many members of the multicast group are connected to the segment. Without native multicast support, routers and switches on shared links would need to use broadcast with software filters, such that every node incurs software overhead for every packet, even if that node is not a member of the multicast group. Alternately, the router would transmit a separate copy to every member of the multicast group on the segment, as is done on multicast incapable switched subnets.

Subnetworks using shared channels, like satellites do, are particularly well suitable for native multicasting, and their designers should make every effort to support it. This involves **designating a section of the subnetwork's own address space for multicasting**. On these networks multicast is basically broadcast on the media, with hardware receiver filters.

Receivers also need to be designed to accept packets addressed to some number of multicast addresses in addition to the unicast packets specifically addressed to them. How many multicast addresses need to be supported by a host depends on the requirements of the associated host; at least several dozen will meet most current needs.

On low-speed networks the multicast address recognition function may be readily implemented in host software, but on high speed networks it should be implemented in subnetwork hardware.

Switched subnetworks must also provide a mechanism for copying multicast packets to ensure the packets reach at least all members of a multicast group. One option is to "flood" multicast packets, in the same manner as broadcast. This can lead to unnecessary transmissions on some subnetwork links, including multiple spotbeam satellite systems. Some subnetworks therefore allow multicast filter tables to control which links receive packets belonging to a specific group. To configure this automatically requires access to layer 3 group membership information.

Multicast group receivers express an interest in receiving a particular data stream where hosts can be located anywhere on the Internet. Hosts that are interested in receiving data flowing to a particular group must join the group using the Internet Group Management Protocol (IGMP). Hosts must be a member of the group to receive the data stream. *Multicast addresses* specify an arbitrary group of IP hosts that have joined the group and want to receive traffic sent to this group. The *Internet Assigned Numbers Authority (IANA)* controls the assignment of IP multicast addresses. IP multicast group addresses will fall in the range of 224.0.0.0 to 239.255.255.255.

Multicasting relies on datagram transmissions, and may use a connection-less protocol such as UDP as its transport protocol, which is basically only a best-effort protocol. Multicast data transport is therefore in principle unreliable, so if specific reliability is required it must be handled at a higher level.

An IP multicast enabled network requires two essential protocol components (IETF RFC 3170):

1) An **IP host-based protocol** to allow a receiver application to notify a local router(s) that it has joined the group, and initiate the data flow from all sender(s) within the scope.

2) An **IP router-based protocol** to allow any routers with multicast group members (receivers) on their local networks to communicate with other routers to ensure that all datagrams sent to the group address are forwarded to all receivers within the intended scope.

Ideally, these protocol components are transparent to multicast applications. However, there are two aspects of their functionality requirements that are worth mentioning specifically, since they affect application performance and design. These are the multicast application requirements for a) Expedient Joins and Leaves and b) Sends without a Join.

4.3 IP broadcasting

Broadcast may be considered a special case of multicast, and in fact in IPv6 it only exists as a part of multicast. So, broadcast is a multicast group whose members include all members in a subnetwork.

Subnetworks can be point-to-point or shared. A point-to-point subnet has exactly two endpoint components (hosts or gateways); a shared link has more than two, using either an inherently broadcast media (e.g. Ethernet, radio, satellite) or on a switching layer hidden from the network layer (switched Ethernet, ATM). Switched subnetworks handle broadcast by copying broadcast packets to ensure each end system receives a copy of each packet.

Several Internet protocols (briefly explained below) for IPv4 make use of broadcast capabilities, including link-layer address lookup (ARP), auto-configuration (RARP, BOOTP, DHCP), and routing (RIP). The lack of broadcast can impede the performance of these protocols, or render them inoperable (e.g. DHCP). ARP-like link address lookup can be provided by a centralized database, but at the expense of potentially higher response latency and the need for nodes to have explicit knowledge of the ARP server address. Shared links should support native, link-layer subnet broadcast. A corresponding set of IPv6 protocols use multicasting instead of broadcast to provide similar functions with improved scaling in large networks.

- Address Resolution Protocol (ARP) is a protocol for mapping an Internet Protocol address (IP address) to a physical machine address that is recognized in the local network.

- RARP (Reverse Address Resolution Protocol) is a protocol by which a physical machine in a local area network can request to learn its IP address from a gateway server's Address Resolution Protocol (ARP) table or cache.

- BOOTP (Bootstrap Protocol) is a protocol that lets a network user be automatically configured (receive an IP address) and have an operating system booted (initiated) without user involvement.

- Dynamic Host Configuration Protocol (DHCP) is a communications protocol that lets network administrators manage centrally and automate the assignment of Internet Protocol (IP) addresses in an organization's network.

- RIP (Routing Information Protocol) is a widely-used protocol for managing router information within a self-contained network such as a corporate Local Area Network (LAN) or an interconnected group of such LANs.

4.4 IP multicast addresses

In IETF RFC 3170, we find mentioned that one of the first questions facing a multicast application developer is what multicast address to use. Multicast addresses are not assigned to individual hosts, assignments can change dynamically, and addresses sometimes have semantics of their own (e.g. Admin Scoping). Multicast applications require an **address management service** that provides address allocation or assignment queries. There are a number of ways for applications to learn about multicast addresses:

- Hard-Coded: Software configuration, encoded in a binary executable, or burned into ROM in embedded devices. These applications typically reference IANA statically allocated multicast addresses (including relative addresses).
- Advertised: Session announcements (as described in the next clause), or via another "out-of-band" query or discovery protocol mechanism.
- Algorithmically Derived: Using a programmatic algorithm to allocate a statistically random (unused) address.

In almost all cases, application designers should assume that multicast addresses are to be dynamic. Very little of the multicast address space is available for static assignment by IANA. Also, given the host-specific addressing available with Source Specific Multicast (SSM), Internet-wide, static address assignment is expected to be very rare.

4.4.1 Class D addresses

Multicast addresses specify an arbitrary group of IP hosts that have joined the group and wish to receive traffic sent to this group.

The Internet Assigned Numbers Authority (IANA) controls the assignment of IP multicast addresses. IANA has assigned the Class D address space to be used for IP multicast. This means that all IP multicast-group addresses will fall in this range: 224.0.0.0 - 239.255.255.255.

The IANA has reserved addresses in the 224.0.0.0 through 224.0.0.255 to be used by network protocols on a local network segment. Packets with these addresses should never be forwarded by a router. They remain local on a particular LAN segment. They are always transmitted with a Time To Live (TTL) of 1. Network protocols use these addresses for automatic router discovery and to communicate important routing information.

The addresses from 224.0.1.0 through 238.255.255.255 are called Globally Scoped Address. They can be used to multicast data between organizations and across the Internet. Some of these addresses have been reserved for use by multicast applications through IANA. For example, 224.0.1.1 has been reserved for Network Time Protocol (NTP).

The addresses from 239.0.0.0 through 239.255.255.255 are called Limited Scope Addresses or Administratively Scoped Addresses. These are defined by IETF RFC 2365 [12] to be constrained to a local group or organization. Routers are typically configured with filters to prevent multicast traffic in this address range from flowing outside of an Autonomous System (AS) or any user defined multicast domain. Within an Autonomous System or domain the Limited Scope address range can be further subdivided so that local multicast boundaries can be defined. This will also allow for address reuse between these smaller domains.

Unlike IPv4 unicast address assignment, where blocks of addresses are delegated to regional registries, IPv4 multicast addresses are assigned directly by the IANA. Current assignments appear as follows:

```
224.0.0.0 - 224.0.0.255 (224.0.0/24) Local Network Control Block
224.0.1.0 - 224.0.1.255 (224.0.1/24) Internetwork Control Block
224.0.2.0 - 224.0.255.0 AD-HOC Block
224.1.0.0 - 224.1.255.255 (224.1/16) ST Multicast Groups
224.2.0.0 - 224.2.255.255 (224.2/16) SDP/SAP Block
224.252.0.0 - 224.255.255.255 DIS Transient Block
225.0.0.0 - 231.255.255.255 RESERVED
232.0.0.0 - 232.255.255.255 (232/8) Source Specific Multicast
Block
233.0.0.0 - 233.255.255.255 (233/8) GLOP Block
234.0.0.0 - 238.255.255.255 RESERVED
239.0.0.0 - 239.255.255.255 (239/8) Administratively Scoped
Block
```

Other important addresses are:

- 224.0.0.4 DVMRP.
- 224.0.0.9 RIP2 Routers.
- 224.0.0.13 PIM Routers.
- 224.0.0.22 IGMPv3 Reports.

More information about reserved multicast addresses can be found here:

http://www.iana.org/assignments/multicast-addresses

4.4.2 MAC layer mapping

To support IP multicasting, the Internet authorities have reserved the multicast address range of 01-00-5E-00-00-00 to 01-00-5E-7F-FF-FF for Ethernet and Fiber Distributed Data Interface (FDDI) media access control (MAC) addresses. The high order 25 bits of the 48-bit MAC address are fixed and the low order 23 bits are variable. To map an IP multicast address to a MAC-layer multicast address, the low order 23 bits of the IP multicast address are mapped directly to the low order 23 bits in the MAC-layer multicast address. Because the first 4 bits of an IP multicast address are fixed according to the class D convention, there are 5 bits in the IP multicast address that do not map to the MAC-layer multicast address. Therefore, it is possible for a host to receive MAC-layer multicast packets for groups to which it does not belong. However, these packets are dropped by IP once the destination IP address is determined.

4.5 IP Multicast Routing Protocols

For multicast, the optimal solution (minimum cost to interconnect N nodes) would impose a (NP-hard) Steiner tree computation. Unfortunately, no multicast routing protocol today is able to maintain such an optimal tree. Different multicast protocols will therefore, in general, generate different trees.

Info:

Steiner tree often arises in network design and wiring layout problems. Suppose we are given a set of sites that must be connected by wires as cheaply as possible. The minimum Steiner tree describes the way to connect them using the smallest amount of wire. Analogous problems arise in designing networks of water pipes or heating ducts in buildings. Similar considerations also arise in VLSI circuit layout, where we seek to connect a set of sites to (say) ground under constraints such as material cost, signal propagation time, or reducing capacitance.

The Steiner tree problem is distinguished from the minimum spanning tree problem in that we are permitted to construct or select intermediate connection points to reduce the cost of the tree.

- Steiner tree definition: A minimum-weight tree connecting a designated set of vertices, called terminals, in a weighted graph or points in a space. The tree may include non-terminals, which are called Steiner vertices or Steiner points.
- NP-hard definition: The complexity class of decision problems that are intrinsically harder than those that can be solved by a nondeterministic Turing machine in polynomial time. When a decision version of a combinatorial optimization problem is proven to belong to the class of NP-complete problems, which includes well-known problems such as satisfiability, travelling salesman, the bin packing problem, etc., then the optimization version is NP-hard.
- Nondeterministic Turing machine definition: A Turing machine which has more than one next state for some combinations of contents of the current cell and current state. An input is accepted if any move sequence leads to acceptance.

All definitions from NIST, the National Institute of Standards and Technology, http://www.nist.gov/.

IP multicast requires destination hosts wanting to receive a multicast to subscribe, using the Internet Group Management Protocol (IGMP) (which supports other related functions, such as leaving a multicast group). Subscribing is done by specifying the Class D IP address used for the particular multicast (IETF RFC 1112 [7]). IGMPv3 adds the ability to specify INCLUDE and EXCLUDE lists based on IP source addresses.

Routers track such IGMP requests and build a connectivity tree for each possible sender to each registered receiver. When multicast traffic is received from a particular sender, the router then uses its tree for that sender to determine on which ports traffic needs to be forwarded.

There are three potential router-to-router protocols to support routers dynamically learning which multicast group's data need to be sent out which ports (that is, the building of the trees):

- Protocol Independent Multicast (PIM), which works with more protocols than just TCP/IP.
- Distance Vector Multicast Routing Protocol (DVMRP, used by MBone).
- Multicast Open Shortest Path First (MOSPF).

A good overview of multicast protocols is given at: http://www-inf.enst.fr/~dax/guides/multicast/protocol.html.

4.5.1 Distance Vector Multicast Routing Protocol (DVMRP)

DVMRP (IETF RFC 1075 [23]) is an interior gateway protocol; suitable for use within an autonomous system, but not between different autonomous systems. DVMRP is not for use in routing non-multicast datagrams, so a router that routes both multicast and unicast datagrams must run two separate routing processes. DVMRP is designed to be easily extensible and could be extended to route unicast datagrams.

DVMRP differs from Routing Information Protocol (RIP) in one very important way. RIP is a simple and now quite old protocol, but has very little overhead in terms of bandwidth used and configuration and management time. RIP is also easy to implement. But RIP thinks in terms of routing and forwarding datagrams to a particular destination, while the purpose of DVMRP is to keep track of the return paths to the source of multicast datagrams. To make explanation of DVMRP more consistent with RIP, the word "destination" is used instead of the more proper "source", but datagrams are not forwarded to these destinations, but originate from them.

4.5.2 Protocol Independent Multicast (PIM)

PIM is a router-to-router protocol that supports multicast traffic over existing unicast routing protocols. PIM was designed to avoid the scaling problems and the potential performance problems. PIM operates in either sparse or dense mode.

There is now (Q1/2003) an Internet draft that defines a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. In particular, it describes managed objects used for managing the Protocol Independent Multicast (PIM) protocol.

4.5.2.1 Dense Mode

Dense Mode (DM) is intended for networks in which most LANs need to receive the multicast (such as LAN TV and corporate and financial information broadcasts). DM uses reverse-path forwarding, where the traffic is initially sent to all routers, and those that do not need a traffic feed will reply with a prune message that removes it from the forwarding list.

Dense-mode routing protocols such as DVMRP and PIM-DM, are not well suited for use over subnetworks with a large round trip delay. Such protocols rely on flooding all multicast packet until they receive an explicit "prune" message.

PIM-DM is usually not considered the best solutions optimal for satellite systems.

4.5.2.2 Sparse Mode

In contrast to Dense Mode, Sparse Mode protocols (e.g. PIM-SM and CBT) do not employ a "Flood and Prune" mode of operation and are much better suited to links with appreciable round trip delay. These protocols are also preferred in an Internet context, and are considered better suited for satellite multicast applications.

Core Based Trees (CBT) is a multicast routing architecture that builds a single delivery tree per group which is shared by all of the group's senders and receivers.

Sparse-mode PIM (PIM-SM) is intended for networks where several different multicasts, often to a small number of receivers, are in progress simultaneously. Senders register their source address with a Rendezvous Point (RP).

Receivers use this (or another) RP to first receive the packets sent to the specified group by the source. Traffic is sent by the sender to the rendezvous point, which then forwards it to the registered receivers. Later an optimal path, if necessary bypassing the rendezvous point, can be created, but traffic is still also sent to the rendezvous point, in anticipation of new receivers registering. Data packets are sent to the RP, but if no receiver registers, the RP will send a "register stop" message and prune the path. The state at the RP is periodically refreshed. If a receiver sends a join to the RP, the RP will join back to the source and traffic will flow via the RP.

PIM Sparse Mode characteristics:

- Routes join a PIM-SM tree by sending explicit join messages.
- Uses Rendezvous Points (RP) for receivers and new sources:
  - Sources join directly to the rendezvous point and send data directly to it.
  - Hosts join toward the rendezvous point via their first-hop routers, building the distribution tree from receiver to source. Unicast routing is employed by the receivers to establish the shortest path.

PIM-SM assumes that nobody on the network wants a packet unless they ask for it via an IGMP report. Rendezvous points are virtual meeting places where senders meet receivers. The rendezvous points advertise themselves on the network and help set up and manages communication between sources and receivers. RPs learn of the address of sources using the Multicast Source Discovery Protocol (MSDP), which describes a mechanism to connect multiple PIM-SM domains together.

When the first hop router receives the multicast packets, it learns the source address, and may choose to send a PIM-SM join directly to each source it finds. When it receives multicast packets along the path from the source, it may prune the shared path via the RP. This may use a more efficient path, and saves RP processing. The RP replicates the packet only if there are downstream group member. Otherwise the RP drops the packets and prunes the source back to the DR which is upstream of it.

When data reaches the rendezvous point, the multicast packets are replicated and sent down the multicast distribution tree toward interested receivers. Replication occurs only at branches of the distribution tree, until packets reach their final destinations when a message is sent back to the first-hop multicast router giving the group address for this service. The ability to replicate information eliminates the need to flood router interfaces with unnecessary traffic.

PIM-SM is used for satellite networks.

4.5.2.2.1 Rendezvous Point (RP)

PIM-SM requires at least one Rendezvous Point, or RP, per domain to function. Initially, receivers do not need to know the location of a source to function, as the address of the RP is distributed throughout the PIM-domain. When a receiver wants to join a group, G, it sends an IGMP member report to its first hop router, which sends a (*,G) join to the RP. In case there is more than one first hop router, one is elected to be the Designated Router (DR), and it carries out this task.

Similarly, when a source wants to begin transmitting to a group, its DR encapsulates and unicast the multicast data to the RP, which strips off the encapsulation and multicasts it to the group members (if any).

The RP is always a router, while a source is a computer attached to a router. Thus, in general a RP is different from the source.

When an RP in a PIM-SM domain first learns of a new sender it constructs a "Source-Active" (SA) message and sends it to its MSDP peers. The SA message contains the source address of the data source, the group address the data source sends to and the IP address of the RP. Each MSDP peer receives and forwards the message away from the RP address in a "peer-RPF flooding" fashion. The Multicast RPF Routing Information Base (MRIB) is examined to determine which peer towards the originating RP of the SA message is selected. Such a peer is called an "RPF peer".

4.6 Inter-domain multicast

4.6.1 General

While intra-domain multicast routing is fairly well established, with Protocol Independent Multicast-Sparse Mode (PIM-SM) accepted as the de facto multicast routing protocol, inter-domain multicast routing presents another set of challenges. The focus is on how services can be shared and distributed between providers. A number of issues must be resolved for this to occur. They include the following:

- Handling differences in topology.
- Handling differences in policy/accounting, etc.
- Avoiding third-party dependencies, i.e. like having to rely on rendezvous points (RP) that lie in another service provider's domain. Sometimes a RP in a different country might be an issue as well.
- Reliable ways for distributed receivers to access distributed sources.
- Handling address resolution across multiple networks and establishing mechanisms for multicast address allocation.

With MBGP, both unicast and multicast routes are carried in the same session but in different routing tables. Because MBGP is an enhanced version of BGP-4, all the familiar policy and configuration tools are available for multicast. The acronym MBGP is often read as Multicast BGP, but the correct name is Multiprotocol BGP.

The Multicast Source Discovery Protocol, MSDP, describes a mechanism to connect multiple PIM-SM domains together. Each PIM-SM domain uses its own independent RP(s) and does not have to depend on RPs in other domains. MSDP has drawbacks associated with join latency, scalability and bursty sources. (MSDP is described in further detail below).

While MBGP and MSDP are steps toward providing interdomain multicast, they alone do not form a complete solution. MBGP is capable of determining the next hop to a host, but not of providing multicast tree construction functions. To answer questions concerning the format of the join message, when join messages should be sent, and how often new inter-domain functionality is needed.

The Border Gateway Multicast Protocol (BGMP) was proposed some time ago as a long-term solution for inter-domain multicast, but it still seems to be at the experimental and academic level. However, the key idea was to construct bi-directional shared trees between domains using a single root. BGMP would then decide in which particular domain to root the shared tree. BGMP assumes that interdomain dependencies can be avoided by using a strict address allocation scheme.

Presently, however, there is a gap here, and new standards are needed. At the moment intra-domain solutions can work fine, but efficient inter-domain solutions need more work and study.

4.6.2 Multicast Source Discovery Protocol (MSDP)

MSDP is found in draft-ietf-msdp-spec-14.txt as of March 2003 [35], and the below text is taken form there.

The Multicast Source Discovery Protocol, MSDP, describes a mechanism to connect multiple PIM-SM domains together. Each PIM-SM domain uses its own independent RP(s) and does not have to depend on RPs in other domains.

Advantages of this approach include:

- No Third-party resource dependencies on RP
  - PIM-SM domains can rely on their own RPs only.
- Receiver only Domains
  - Domains with only receivers get data without globally advertising group membership.

MSDP may be used with protocols other than PIM-SM, but such usage is not specified in this memo.

MSDP-speaking routers in a PIM-SM (IETF RFC 2362 [22]) domain have a MSDP peering relationship with MSDP peers in another domain. The peering relationship is made up of a TCP connection in which control information is exchanged. Each domain has one or more connections to this virtual topology. The purpose of this topology is to allow domains to discover multicast sources from other domains. If the multicast sources are of interest to a domain which has receivers, the normal source-tree building mechanism in PIM-SM will be used to deliver multicast data over an inter-domain distribution tree.

When an RP in a PIM-SM domain first learns of a new sender, e.g. via PIM register messages, it constructs a "Source-Active" (SA) message and sends it to its MSDP peers. The SA message contains the following fields:

- Source address of the data source.
- Group address the data source sends to.
- IP address of the RP.

Each MSDP peer receives and forwards the message away from the RP address in a "peer-RPF flooding" fashion. The notion of peer-RPF flooding is with respect to forwarding SA messages. The Multicast RPF Routing Information Base (MRIB) is examined to determine which peer towards the originating RP of the SA message is selected. Such a peer is called an "RPF peer".

The procedure outlined in MSDP has been named flood-and-join because if any RP is not interested in the group, they can ignore the SA message. Otherwise, they join a distribution tree.

A MSDP speaker MUST cache SA messages. Caching allows pacing of MSDP messages as well as reducing join latency for new receivers of a group G at an originating RP which has existing MSDP (S,G) state. Caching also aids in diagnosis and debugging of various problems. The main timers for MSDP are: SA-Advertisement-Timer, SA Cache Entry timer, Peer Hold Timer, KeepAlive timer, and ConnectRetry timer.

The key advantage of MBGP/PIM-SM/MSDP is that it is a functional solution largely built on existing protocols. Furthermore, it is already being deployed with a fair amount of success. The key disadvantage is that, as a long-term solution, the MBGP/PIM-SM/MSDP protocol suite may be susceptible to scalability problems.

4.7 Internet Group Management

4.7.1 IGMP

The Internet Group Management Protocol (IGMP) is used to provide information about membership of multicast groups on a network. Hosts that wish to join a multicast group instruct their network adapter cards to listen for network frames specifying a special MAC address that corresponds to the multicast class D IPv4 address of the group. In practice, the mapping is not unique, so the IPv4 module also has to perform some filtering of received frames.

Multicast hosts must inform their nearest routers that they should receive multicast messages for selected multicast groups. The Internet Group Management Protocol (IGMP) is an Internet protocol that provides a way for an Internet computer to report its multicast group membership to adjacent routers. IGMP is also used by routers to regularly verify whether known group members are still active. On the basis of IGMP information routers will decide if multicast messages it receives shall be forwarded or not. When receiving a multicast packet, the router will check if there is at least one group member its subnetwork, and if so the router will forward the message.

If there are several multicast routers on a subnetwork, one of the routers must be responsible for keeping the membership state of the multicast groups on the subnetwork.

IETF RFC 1112 [7] defines the specification for IGMP Version 1. In Version 1, there are two different types of IGMP messages:

- Membership Query.
- Membership Report.

Hosts send out IGMP Membership Reports for a particular multicast group to join. The router periodically sends out an IGMP Membership Query to verify that at least one host on the subnet is still active in receiving traffic for a given group. If there are no replies to three consecutive IGMP Membership Queries, a router will timeout the group and stop forwarding traffic directed toward that group.

IGMPv1 is not recommended for use anymore.

IETF RFC 2236 [8] defines the specification for IGMP Version 2. In Version 2, there are four types of IGMP messages:

- Membership Query.
- Version 1 Membership Report.
- Version 2 Membership Report.
- Leave Group.

The main difference from version 1 is that there is a Leave Group message, where hosts actively notify the local multicast router if they want to leave a group. The router then sends out a group specific query and to see if there are any remaining hosts interested the group. If not, the router will stop forwarding that traffic. This can reduce the leave latency compared to Version 1, and unnecessary traffic can be stopped sooner.

Version 3 of IGMP (IETF RFC 3376 [24]) adds support for "source filtering", that is, the ability for a system to report interest in receiving packets only from specific source addresses, or from all but specific source addresses, sent to a particular multicast address. That information may be used by multicast routing protocols to avoid delivering multicast packets from specific sources to networks where there are no interested receivers.

IGMP Version 3 makes it possible for a host to join a group and specify a set of sources of that group from which it wants to receive multicast messages. Similarly, leave group messages of Version 2 has been enhanced to support group-source leave messages. Membership Queries are sent by IP multicast routers to query the multicast reception state of neighbouring interfaces.

IGMP version 3, allows for specific joins and leaves, through the addition of source specific INCLUDE reports, so that it will be possible to join a specific source of a specific group directly. This capability will make Source Specific Multicasting (SSM) possible. Windows XP supports IGMPv3 natively.

When routers are running different versions of IGMP, the routers negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

It is possible to create IGMP static group membership, for instance to test multicast forwarding without a receiver host. When IGMP static group membership is enabled, data is forwarded to an interface without receiving membership reports from host members.

There are a number of tunable parameters and timers in IGMP. Most of these timers are configurable. If non-default settings are used, they MUST be consistent among all systems on a single link. Some of the most relevant are listed below:

- The Query Interval is the interval between General Queries sent by the Querier. Default: 125 s. By varying the [Query Interval], an administrator may tune the number of IGMP messages on the network; larger values cause IGMP Queries to be sent less often.

    - The Max Response Time used to calculate the Max Resp Code inserted into the periodic General Queries. Default: 100 (10 s). By varying the [Query Response Interval], an administrator may tune the burstiness of IGMP messages on the network; larger values make the traffic less bursty, as host responses are spread out over a larger interval.

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

- The overall level of periodic IGMP traffic is inversely proportional to the Query Interval. A longer Query Interval results in a lower overall level of IGMP traffic. The Query Interval MUST be equal to or longer than the Max Response Time inserted in General Query messages.

- The Last Member Query Interval is the Max Response Time used to calculate the Max Resp Code inserted into Group-Specific Queries sent in response to Leave Group messages. Default: 10 (1 s). This value may be tuned to modify the "leave latency" of the network. A reduced value results in reduced time to detect the loss of the last member of a group or source.

- The Robustness Variable tunes IGMP to expected losses on a link. IGMPv3 is robust to (Robustness Variable - 1) packet losses, e.g. if the Robustness Variable is set to the default value of 2, IGMPv3 is robust to a single packet loss but may operate imperfectly if more losses occur. On lossy subnetworks, the Robustness Variable should be increased to allow for the expected level of packet loss. However, increasing the Robustness Variable increases the leave latency of the subnetwork. (The leave latency is the time between when the last member stops listening to a source or group and when the traffic stops flowing).

- The burstiness of IGMP traffic is inversely proportional to the Max Response Time. A longer Max Response Time will spread Report messages over a longer interval. However, a longer Max Response Time in Group-Specific and Source-and-Group-Specific Queries extends the leave latency. (The leave latency is the time between when the last member stops listening to a source or group and when the traffic stops flowing). The expected rate of Report messages can be calculated by dividing the expected number of Reporters by the Max Response Time. The Max Response Time may be dynamically calculated per Query by using the expected number of Reporters for that Query as follows.

Query Type	Expected number of Reporters
General Query	All systems on subnetwork
Group-Specific Query	Systems that had expressed interest
Source-and-Group- Specific Query	All systems on the subnetwork that had expressed interest in the source and group

GMPv2 suffers some scaling issues similar to, but different from, those encountered with dense-mode routing protocols:

- A large subnetwork round trip transit delay increases reduces the effectiveness of the feedback suppression technique. The number of membership reports received per membership query is increased when there are many members of the same set of groups.

- Too small a query interval may result in multicast routers assuming that there are no responders, when membership report messages are actually still in flight over the subnetwork. This case could lead to interruption of the multicast service.

- A large query interval may lead to a significant delay in removing unnecessary multicast traffic.

IGMPv3 does not use a suppression technique to reduce the number of membership reports. The volume of IGMPv3 traffic is not impacted by the subnetwork round-trip transit delay. A small query interval may lead to interruption of the service, but the "leave processing" in IGMPv3 is not normally impacted by query interval.

4.7.2 PIM-SM/IGMPv3 interaction

A PIM-SM/IGMPv3 interaction takes place when a PM-SM router receives an IGMP message regarding a group address that is in the Any Source Multicast (ASM) range. This range is defined as the entire Class D multicast space excluding the global SSM range and any locally defined Source Specific space.

PIM-SM join messages are initiated when a PIM-SM router determines that there are entities interested in a specific group or a specific source sending to the group. If this is due to a IGMPv3 report with a zero-length EXCLUDE list, then the join is sent as a (*,G) join towards the RP.

If the join is triggered by the reception of an IGMPv3 report that contains source specific information, the join is sent as a (S,G) join towards the specific source. This behaviour optimizes the join process, as well as facilitates the adoption of the SSM model. It also can cause failures in some specific network architectures, and thus, can be overridden by local policy. If this is the case, then all IGMPv3 triggered joins are sent towards the RP as (*,G) joins. The initiating router is responsible for filtering the data before forwarding to the requesting network.

PIM-SM prune messages are initiated when a PIM-SM router determines that there are no entities interested in a specific group, or a specific source sending to the group. If this is triggered by either receiving an IGMP report with an EXCLUDE or if a specific IGMP derived Source/Group times out, then an (S,G) prune is sent towards the upstream router. If all of the IGMP derived requests for a group time out, then (S,G) and (*,G) prunes are sent upstream as needed to stop all flow of traffic for that group.

4.8 Source Specific Multicast

Source Specific Multicasting (SSM) involves using the capabilities of IGMP v3 to tune PIM-SM to the needs of large scale, one-to-many, multi-casting, such as in web casting. IGMP v3 allows a source to explicitly request traffic from a particular (S,G) pair without using a Rendezvous Point at all. Edge routers must be somewhat modified but routers in the interior of the network do not, and could run standard PIM-SM. Knowledge of the source and group pair is assumed to be known a priori, such as from a web-page.

SSM is included as part of the PIM-SM v. 2 (draft). The address range 232/8 has been reserved for SSM.

From the Internet draft "An Overview of SSM Deployment" (4 December 2001) [36], we find that the Source Specific Multicast (SSM) service model defines a "channel" identified by an (S,G) pair, where S is a source address and G is an SSM destination address. Channel subscriptions are described using an SFM-capable group management protocol such as IGMPv3 or MLDv2. Only source-based forwarding trees are needed to implement this model.

The SSM service model alleviates all of the deployment problems of some present solutions by:

- Address Allocation: SSM defines channels on a per-source basis, i.e. the channel (S1,G) is distinct from the channel (S2,G), where S1 and S2 are source addresses, and G is an SSM destination address. This averts the problem of global allocation of SSM destination addresses, and makes each source independently responsible for resolving address collisions for the various channels that it creates.

- Access Control: SSM lends itself to an elegant solution to the access control problem. When a receiver subscribes to an (S,G) channel, it receives data sent by an only the source S. In contrast, any host can transmit to an ASM host group. At the same time, when a sender picks a channel (S,G) to transmit on, it is automatically ensured that no other sender will be transmitting on the same channel (except in the case of malicious acts such as address spoofing). This makes it much harder to "spam" an SSM channel than an ASM multicast group.

- Handling of well-known sources: SSM requires only source-based forwarding trees; this eliminates the need for a shared tree infrastructure. In terms of the IGMP/PIM-SM/MSDP/MBGP protocol suite, this implies that neither the RP-based shared tree infrastructure of PIM-SM nor the MSDP protocol is required. Thus the complexity of the multicast routing infrastructure for SSM is low, making it viable for immediate deployment. Note that MBGP is still required for distribution of multicast reachability information.

- It is widely held that point-to-multipoint applications such as Internet TV will be important in the near future. The SSM model is ideally suited for such applications.

4.9 Time To Live (TTL) for multicast packets

The IP multicast routing protocol uses the **Time to Live (TTL)** field of IP datagrams to decide how far from a sending host a given multicast packet should be forwarded. The default TTL for multicast datagrams is 1, which will result in multicast packets going only to other hosts on the local network. A **setsockopt**(2) call may be used to change the TTL.

As the values of the TTL field increase, routers will expand the number of hops they will forward a multicast packet. To provide meaningful scope control, multicast routers enforce the following "thresholds" on forwarding based on the TTL field:

0 restricted to the same host;

1 restricted to the same subnet;

32 restricted to the same site;

64 restricted to the same region;

128 restricted to the same continent;

255 unrestricted.

4.10 Multicast scoping

Not all groups need or should have a global scope. Scoping can be used to limit the region in which data is forwarded, which is useful for performance reasons with flood and prune multicast routing protocols. It also useful for application security reasons or because multicast addresses are a scarce resource. Scoping a session may allow the same multicast address to be used several non-overlapping places.

Multicast scoping can be performed in two ways which are known as TTL scoping and administrative scoping, where TTL scoping is most common.

4.10.1 TTL scoping

When an IP packet is sent, the IP header field Time to Live (TTL) is set to a value between zero and 255. Every time a router forwards the packet, it decrements the TTL field in the packet header, and if the value reaches zero, the packet is dropped. (TTL should in principle also be decremented if a packet is queued for more than a certain amount of time).

TTL is normally set to a fixed value in unicast by the sending host (64 and 255 are commonly used) and is intended to prevent packets looping forever, and also forms a part of the formal proof that the TCP close semantics are safe.

TTL can be used to constrain how far a multicast packet can travel by choosing the value put into packets as they are sent. The TTL value can be decremented by more than one to ensure packets do not escape the subnet in question.

4.10.2 Administrative scoping

Administrative scoping is much more flexible than TTL scoping, but suffers from a number of disadvantages. In particular, it is not possible to tell from the address of a packet where it will go unless all the scope zones that the sender is within are known. Also, as administrative boundaries are bi-directional, one scope zone nested within or overlapping another must have totally separate address ranges. This makes their allocation difficult from an administrative point of view, as the ranges ought to be allocated on a top-down basis (largest zone first) in a network where there is no appropriate top-level allocation authority. Also, it is easy to discomfiture a boundary by omitting or incorrectly configuring one of the routers - with TTL scoping it is likely that in many cases a more distant threshold will perform a similar task lessening the consequences, but with administrative scoping there is less likelihood that this is the case.

Scoped addresses also have a number of advantages - particularly since they align to multicast domains, and therefore allow a high assurance that packets intended for only a specific domain do not leak out to other domains by mistake, which can be a big problem for TTL-based schemes.

4.11 Reliable multicast protocols and architectures

The following classifications are based on [21] (M.Koyabe and G. Fairhurst. "Reliable Multicast Via Satellite: A Comparison Survey and Taxonomy", International Journal for Satellite Communications (IJSC), Vol:24(1), 21-26, 2001).

4.11.1 One-to-many (star-based)

One-to-many (Star-based) reliable multicast protocols, Reliable Multicast Transport Protocols are suited for data delivery from a single sender to multiple receivers, with minimal dependency on network elements. Most protocols use a receiver- or sender- initiation approach to achieve reliability.

In a number of star-based multicast protocols, receivers send unicast packet to request for retransmissions, which are later multicast back to requesting receivers. Most use rate-based flow control with random or probabilistic timers to suppress feedback implosion to sender.

The decision about who participates (e.g. responds to feedback packets) in a multicast group is known as "locus control". It may be centralized or distributed, usually group membership is implicit. One-to-many (star-based) protocols support both bulk delivery and streaming. They trade low latency for higher scalability. MFTP makes transfer blocks as large as possible such that, one NACK can represent thousands or tens of thousand packets. This reduces NACK implosion to the sender and achieves better scaling compared to most one-to-many (star-Based) reliable multicast protocols.

The ability to detect and repair losses using Forward Error Correction (FEC) has gained acceptance as a solution offering both reliability and scalability. Many reliable multicast protocols support FEC for various uses. Use of FEC layering with RMDP has been proposed, but is yet to be deployed extensively.

Examples of one-to-many (star-based) reliable multicast protocols are listed below:

- Reliable Multicast Transport Protocol (RMTP);
- Multicast Dissemination protocol (MDP/MDPv2);
- Reliable Multicast data Distribution Protocol (RMDP);
- Multicast Transport Protocol (MTP);
- Multicast File Transfer Protocol (MFTP);
- Adaptive File Distribution Protocol (AFDP);
- Restricted Reliable Multicast Protocol (RRMP).

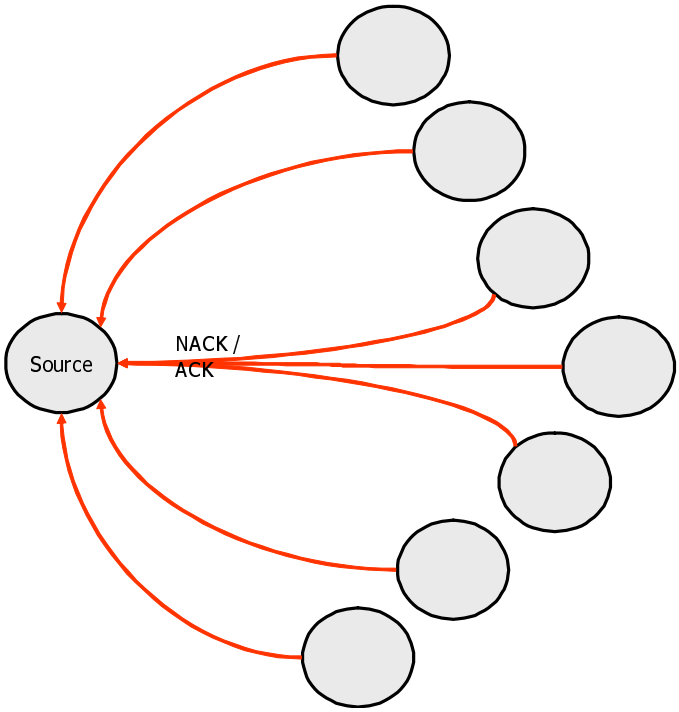

Figure 2: One-to-many, star

4.11.2 One-to-many (tree-based)

One-to-many (tree-based) protocols are reliable multicast protocols, which aggregate receivers into a tree-based structure. A set of multicast groups are used to establish a hierarchy of logical areas (one address per area) and retransmission responsibility is distributed over the acknowledgement tree (ACK tree) structure with the source as the root of the tree. The tree structure prevents receivers from directly contacting the source to maintain scalability of feedback and repair packets when used with a large number of receivers.

Tree-based reliable multicast protocols achieve reliability using either a receiver- or sender- initiated approach. Most one-to-many protocols use a structured approach based on a designated server or parent node, to suppress feedback implosion to the sender. Tree-based Multicast Transport Protocol (TMTP) however, adopts a different approach by organizing multicast members into defined structures in order to filter the amount of feedback generated by the group. Like Star-based, Tree-based protocols use a rate-based approach to regulate flow control.

Due to delays introduced at each layer, tree-based approaches are less favoured for time critical data distribution. Tree-based protocols aggregate receivers in tree-structures, therefore they can only offer improved scaling for terrestrial networks with many intermediate routers, but do not improve performance for satellite networks where the majority of receivers may be reached via a single satellite hop.

Examples of one-to-many (tree-based) protocols are listed below:

- Reliable Multicast Transport Protocol (RMTP/RMTP+);
- Reliable Adaptive Multicast Protocol (RAMP);
- Log-Based Receiver Reliable Multicast (LBRM);
- Local Group based Multicast Protocol (LGMP);
- Tree-based Multicast Transport Protocol (TMPT);
- On-Tree Efficient Recovery using Subcasting (OTERS);
- Tree-based Reliable Multicast (TRAM).

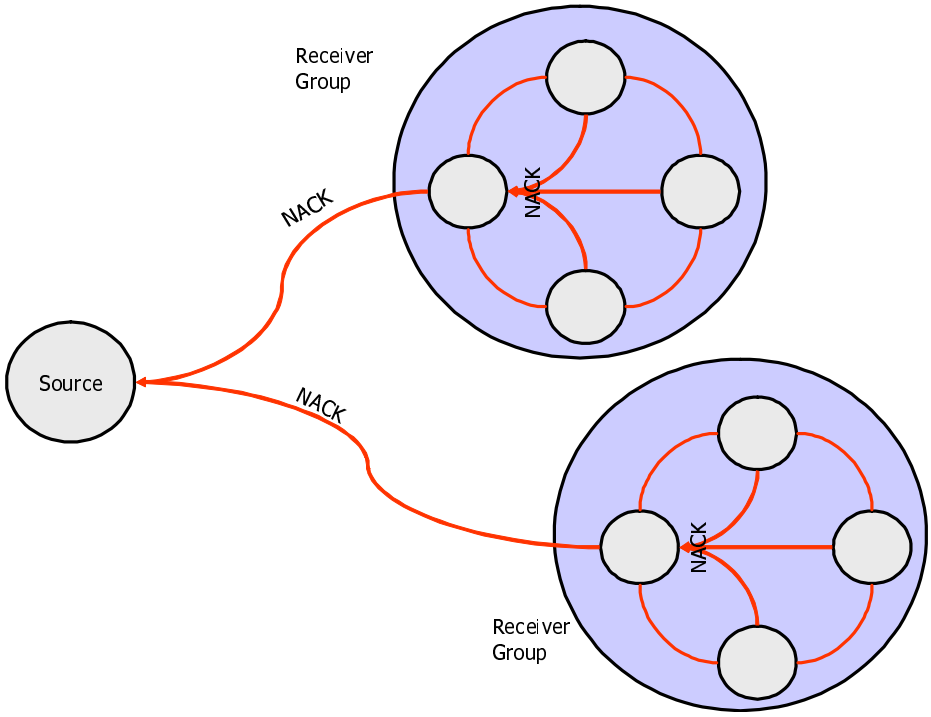

Figure 3: Tree based group organization

4.11.3 One-to-many (ring-based)

Ring-based protocols, Reliable Multicast Protocol (RMP) organize nodes into a ring. Each receiver in a ring maintains global session membership information. A token is based among all receivers to synchronies data transmissions and feedback of acknowledgements. It may also be used to synchronies the order of transmission when multiple senders transmit to a common group of receivers.

RMP uses a receiver- and sender- initiated hybrid approach for reliability and rate-based flow control that allows the administrator to specify the maximum amount of traffic that any sender can send in a given time period. Ring-based topologies have the advantage of high performance and high reliability. However, the complexity of the algorithm and the large number of protocol states make them difficult to implement. Since each member must maintain a membership list, like tree-based protocols, the scalability of ring-based protocols is limited. They are not therefore well suited to wide-area satellite multicast.

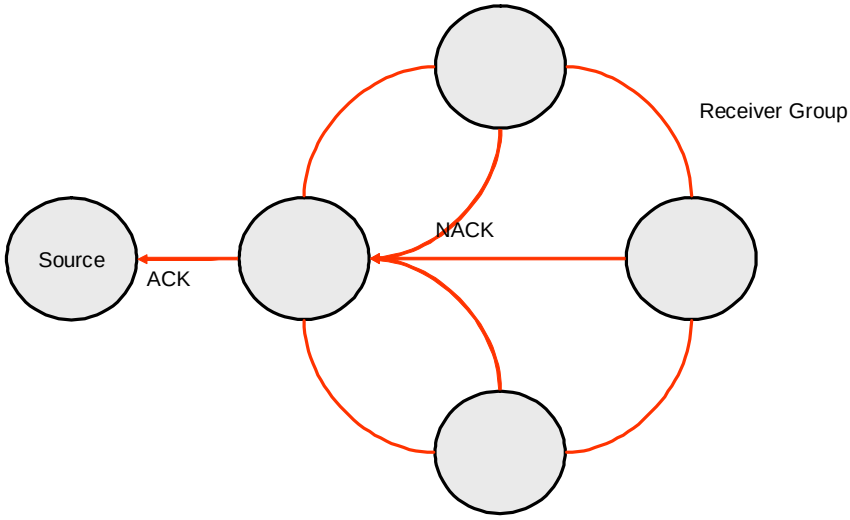

Figure 4: One to many, ring based

4.11.4 Many-to-many

The following are many-to-many protocol designs:

- Scalable Reliable Multicast (SRM);
- Pragmatic General Multicast (PGM);
- Transport Protocol for Reliable Multicast (TRM);
- Reliable Concurrent Multicast from Bursty Sources (RCMBS).

Most many-to-many protocol designs are based on the SRM protocol. Both SRM and TRM support streaming and collaborative applications, while PGM supports both streaming and bulk transfer.

This work will not consider many-to-many multicasting further.

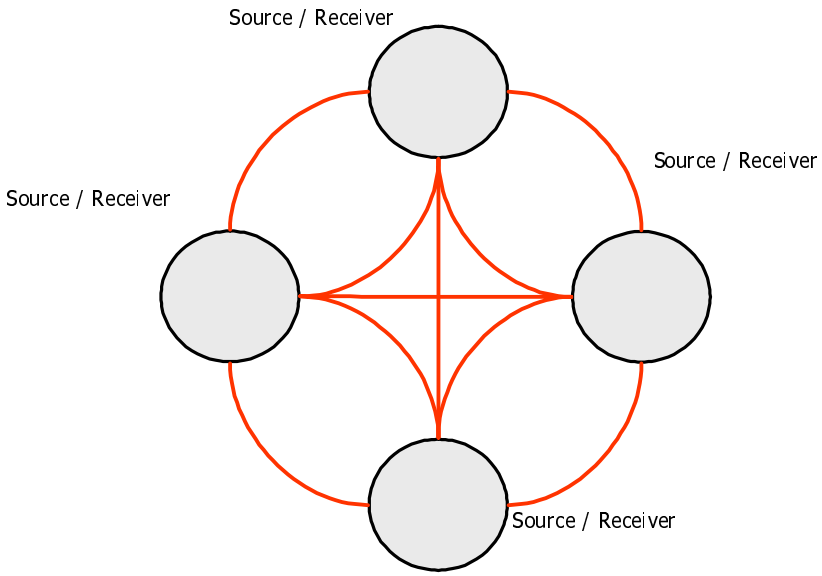

Figure 5: Many to many

4.12 Multicast security

If the multicast transmission is of a sensitive nature, it should be encrypted. This means that the all members of the group must share the same encryption key to take benefit of the multicast transmission.

There are many factors that must be taken into account when developing the desired key management architecture. Important issues for key management architectures include level (strength) of security, cost, initializing the system, policy concerns, access control procedures, performance requirements and support mechanisms. See for instance IETF RFC 2627 [15]. This report contains a discussion of the difficult problem of key management for multicast communication sessions. This RFC focuses on two main areas of concern with respect to key management, which are, initializing the multicast group with a common net key and rekeying the multicast group.

One way of setting up a group of symmetric keys is with the assistance of a centralized key management facility. This facility would act as a key broker creating a distributing key to qualified group members. There are several problems with this centralized concept.

One of the problems with a centralized key distribution system is the concentration of key management workload at a single site. Another drawback is the use of a central site that would process a request in accordance with its priority and current workload. Latencies could develop. A centralized key distribution site is a valuable target for someone to compromise, and expensive procedural security mechanisms are sometimes needed.

Delegating key management responsibility to the groups eliminates the centralized key management site as a single point of failure. The Group Key Management Protocol delegates the access control, key generation, and distribution functions to the communicating entities themselves rather than relying on a central function.

Group Key Management Protocol (GKMP) Specification (IETF RFC 2093 [25]) and the Group Key Management Protocol (GKMP) Architecture (IETF RFC 2094 [26]) propose a protocol to create grouped symmetric keys and distribute them amongst communicating peers. This protocol has the following advantages: 1) virtually invisible to operator, 2) no central key distribution site is needed, 3) only group members have the key, 4) sender or receiver oriented operation, 5) can make use of multicast communications protocols.

Scalable Multicast Key Distribution (experimental IETF RFC 1949 [27]) proposes a protocol to create grouped symmetric keys and distribute them amongst communicating peers. This protocol has the following advantages: 1) virtually invisible to operator, 2) no central key distribution site is needed, 3) only group members have the key, 4) sender or receiver oriented operation, 5) can make use of multicast communications protocols.

4.13 IPv6 and IPv4 multicast issues

Multicast applications have been developed for both IPv4 and IPv6. However IPv6 extends the IP multicasting capabilities by defining a much larger multicast address space. All of the IPv6 hosts and routers are required to support multicast. IPv6 also has no broadcast address as such; it obtains various multicast addresses of various scopes. Improved scope in IPv6 promises to simplify the use and administration of multicast in many applications.

IP v6 issues related to multicasting include features such as:

- greater addressing space;
- better routing performance and services.

Multicasting is widely used in IPv6 which does not support broadcasting at all. IPv6 implements group management within the ICMPv6 protocol. Since PIM is designed to support multiple protocols including IPv6, an existing implementation of PIM for IPv4 can easily be made to accommodate IPv6. With IPv6, it is probable that all routers will support IPv6 multicast, so there may not be a need to use tunnels. If so the multicast and unicast topologies will be the same.

IETF RFC 2375 discusses addressing. IPv6 allows three types of addresses:

- Unicast: An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.
- Anycast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" according to the routing protocols' measure of distance).

- Multicast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address.

Anycast addresses

An anycast address enables a source to specify that it wants to contact any one node from a group of nodes via a single address. A packet with such an address will be routed to the nearest interface in the group, according to the router's measure of distance. An example of the use of an anycast address is within a routing header to specify an intermediate address along a route. The anycast address could refer to the group of routers associated with a particular provider or particular subnet, thus dictating that the packet be routed through that provider or internet in the most efficient manner.

Anycast addresses are allocated from the same address space as unicast addresses. Thus, members of an anycast group must be configured to recognize that address, and routers must be configured to be able to map an anycast address to a group of unicast interface addresses.

One particular form of anycast address, the subnet-router anycast address, is predefined. The subnet prefix field identifies a specific subnetwork. For example, in a provider-based global address space, the subnet prefix is of the form (010 + registry ID + provider ID + subscriber ID + subnet ID). Thus, the anycast address is identical to a unicast address for an interface on this subnetwork, with the interface ID portion set to zero. Any packet sent to this address will be delivered to one router on the subnetwork; all that is required is to insert the correct interface ID into the anycast address to form the unicast destination address.

Multicast Addresses

IPv6 includes the capability to address a predefined group of interfaces with a single multicast address. A packet with a multicast address is to be delivered to all members of the group. A multicast address consists of an 8-bit format prefix of all ones, a 4-bit flags field, a 4-bit scope field, and a 112-bit group ID.

IETF RFC 3307 [16] specifies guidelines that MUST be implemented by any entity responsible for allocating IPv6 multicast addresses. This includes, but is not limited to, any documents or entities wishing to assign permanent IPv6 multicast addresses, allocate dynamic IPv6 multicast addresses, and define permanent IPv6 multicast group identifiers. The purpose of these guidelines is to reduce the probability of IPv6 multicast address collision, not only at the IPv6 layer, but also at the link-layer of media that encode portions of the IP layer address into the link-layer address.

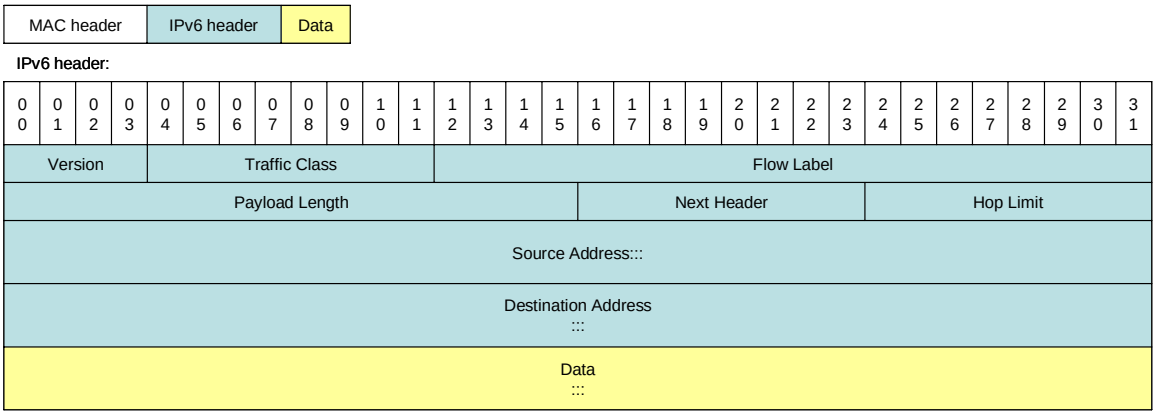

Figure 6: IPv6 header

4.14 Streaming and streaming protocols

4.14.1 Real Time Streaming Protocol (RTSP)

The Real Time Streaming Protocol (RTSP) is an application-level protocol for control over the delivery of data with real-time properties. RTSP provides an extensible framework for controlled, on-demand delivery of real-time data, which can be both live data feeds and stored clips. RTSP is intended to control multiple data delivery sessions, provide a means for choosing delivery channels such as UDP, multicast UDP and TCP, and provide a means for choosing delivery mechanisms based upon RTP (IETF RFC 1889 [11]).

The Real-Time Streaming Protocol establishes and controls either a single or several time-synchronized streams of continuous media such as audio and video. RTSP acts as a "network remote control" for multimedia servers. It does not typically deliver the continuous streams itself, although interleaving of the continuous media stream with the control stream is possible. The set of streams to be controlled is defined by a presentation description.

There is no notion of an RTSP connection; instead, a server maintains a session labelled by an identifier. An RTSP session is in no way tied to a transport-level connection such as a TCP connection. During an RTSP session, an RTSP client may open and close many reliable transport connections to the server to issue RTSP requests. Alternatively, it may use a connectionless transport protocol such as UDP.

4.14.2 The Real-time Transport Protocol (RTP)

IETF RFC 1889 [11] specifies the Real-time Transport Protocol (RTP), which provides end-to-end delivery services for data with real-time characteristics, such as interactive audio and video. Those services include payload type identification, sequence numbering, time-stamping and delivery monitoring. Applications typically run RTP on top of UDP to make use of its multiplexing and checksum services; both protocols contribute parts of the transport protocol functionality. However, RTP may be used with other suitable underlying network or transport protocols. RTP supports data transfer to multiple destinations using multicast distribution if provided by the underlying network.

RTP itself does not provide any mechanism to ensure timely delivery or provide other quality-of-service guarantees, but relies on lower-layer services to do so. It does not guarantee delivery or prevent out-of-order delivery, nor does it assume that the underlying network is reliable and delivers packets in sequence. The sequence numbers included in RTP allow the receiver to reconstruct the sender's packet sequence, but sequence numbers might also be used to determine the proper location of a packet, for example in video decoding, without necessarily decoding packets in sequence. RTP is primarily designed to satisfy the needs of multi-participant multimedia conferences.

IETF RFC 1889 [11] consists of:

- the Real-time Transport Protocol (RTP); and
- the RTP Control Protocol (RTCP), to monitor the quality of service and to convey information about the participants in an on-going session.

5 Applications and use cases

5.1 Introduction

Multicast applications such as caching, streaming, and reliable transfer are not specific to satellite - but there may be satellite specific design considerations such as placement of caches (and other multicast agents), the effect of delay, the effect of subnetwork outages, and potential interactions with bandwidth-on-demand techniques. These may influence the design and configuration of systems intending to use broadband satellite networks. There are implications on using long delay subnetworks.

There here is also a question of scale. Most current multicast applications in use on the Internet address only a modest (less than a hundred) number of receivers. Future deployment will require scaling to larger groups, and satellite subnetworks may be the first place to see the need for multicast groups with a large number of receivers.

We will not try to list all possible satellite multicast applications here, but one of the most prominent applications for multicast may be streaming. Multicast streaming could lend itself to efficient satellite transmission. Streaming media is moving pictures and/or sound sent, usually in compressed form (over the Internet), and presented at the destination (usually) soon after they arrive. The user therefore does not have to download a large file before seeing the video or hearing the sound, since the media is sent in a continuous stream and can be played as it arrives. Streaming video can be sent either from prerecorded video files or be distributed as part of a live broadcast feed.

A multicast stream may also be cached for later viewing. Caching technology can allow larger multicast groups to benefit from the same stream, thus reducing the average amount of resources required to transfer a stream to an end user.

The services in question would use the BSM bearer services to provide BSM multicast, as indicated below.

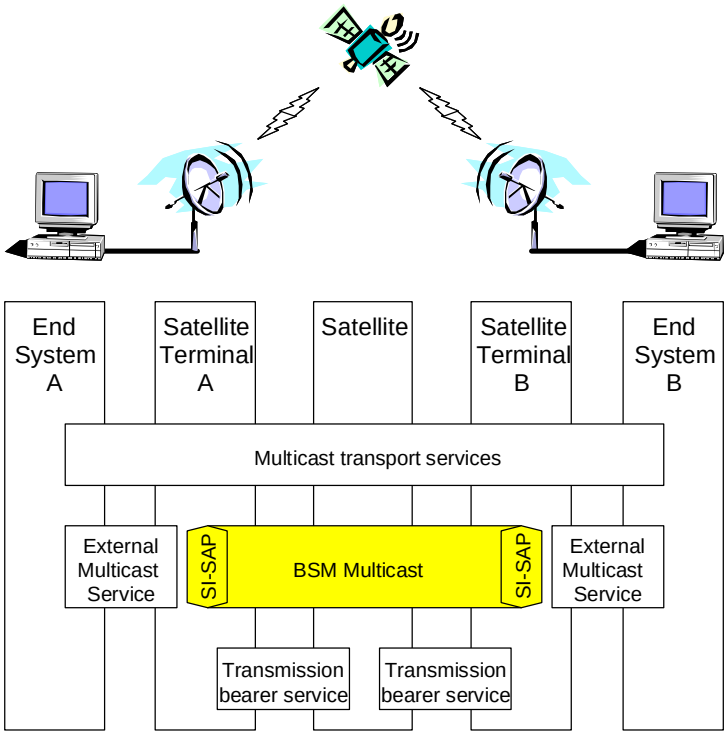

Figure 7: BSM services

5.2 Multicast Backbone (MBone)

The Multicast Backbone, MBone was created in 1992 as an interconnected set of subnetworks with routers capable of forwarding multicast packets for experimenting with multicasting. The MBone consists of servers equipped to handle the multicast protocol. An MBone router that is sending a packet to another MBone router through a non-MBone part of the network encapsulates the multicast packet as a unicast packet. The non-MBone routers simply see an ordinary packet. With an increase of multicast routing software, this has gradually replaced the use of tunnels.

Most of the MBone is now historical. Many sites previously on the MBone have now abandoned their unix-based DVMRP tunnels for routers implementing PIM-SM. Connectivity is common via native IP links or via ISPs offering commercial multicast services.

5.3 Use scenarios

5.3.1 Framework

Multicast over satellite is not limited to broadband Ka/Ku - band satellite systems. Multicast services can also be delivered over L/S/C bands, as well as V/Q band in the future. However, multicast solutions favour the always-connected type of systems.

Only a portion of all broadband Internet connections could be via satellite. Of the ones on satellite, there will be options for one of several different satellite systems. With respect to multicast many hosts can be outside any given satellite system.

Only some multicast hosts (A) are on a satellite network, as illustrated below. Most IP multicast hosts are not expected to be on the satellite network in question. There may be more than one satellite system (B). There may also be hosts that only operate within a given satellite network (C).

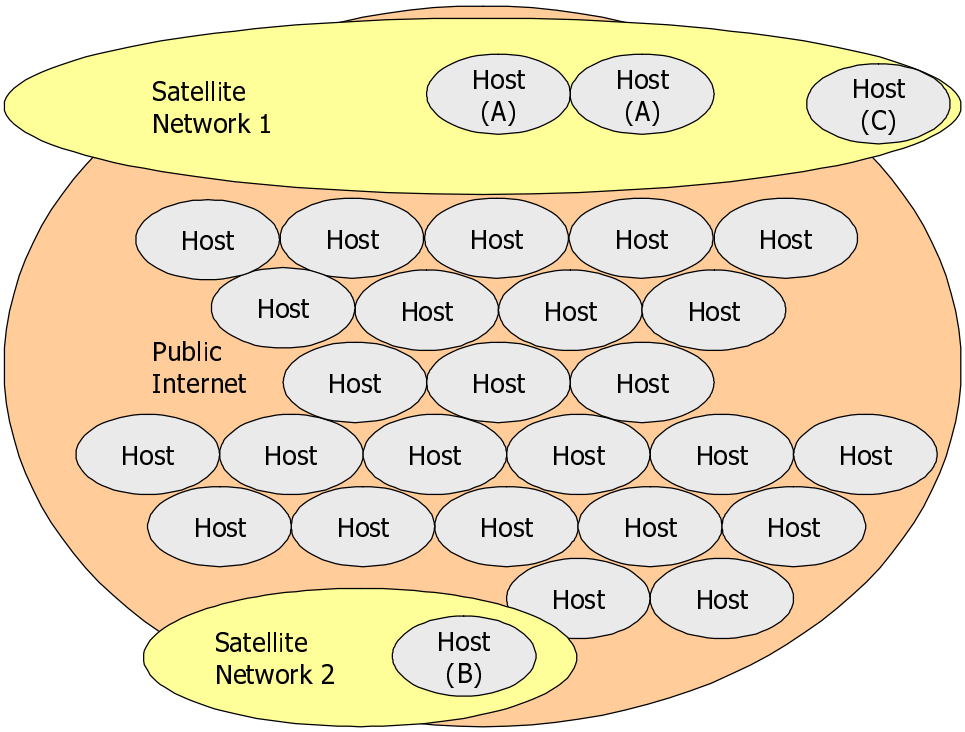

Figure 8: Multicast host distribution example

5.3.2 Open and closed groups

The following use scenarios are considered:

- Open multicast group that can be joined from any Internet connected terminal, including satellite connections.
  - Caching of multicast content may be done at the gateways or at the terminals.
- Closed multicast groups that are limited to the satellite network in question.
  - Virtual private networks are included.
  - Satellite service providers may offer specific content for multicast on the satellite network.
- There may or may not be return channels.
- There may or may not be caching involved, either at the gateway, locally at the terminal or both places.

With open multicast groups it is meant publicly available groups, also visible outside of the satellite system. Closed groups are available only at the specific network in question. Note that not all multicast groups that technically may be available may in fact be offered to any user.

5.3.3 Categories

Two main categories for broadband services have been specified from the network perspective: interactive services and distribution services. The interactive services are subdivided into three classes of services:

- Conversational services: Typical examples are video telephony, video-conferences, video/audio information transmission, high speed digital information, file and document transfer;
- Message services: Typical examples are video mail and document mail; and
- Retrieval services: Typical examples are video, high-resolution image, document and data.

The distribution services are subdivided into two classes:

- the class without user individual representation control, such as TV, multimedia video and audio distribution; and
- the class with user individual representation control, such as Pay TV (PTV).

From the user's perspective, satellite IP multicast should support existing multicast applications such as reliable file transfer, data distribution and multimedia streaming including video, voice and data.

All these services may have different QoS requirements such as jitter sensitive real time or loss sensitive transaction data.

The value added services provided by satellites include extended coverage and efficient delivery to users on a large scale.

Examples of One-to-Many Multicast Services:

- Internet TV, Web casting.
- Web casting of Broadband Streaming Media.
- Remote Education.
- Distribution of Financial Data: Stock-ticker.

And Many-to-Many Multicast:

- Teleconferencing.
- Whiteboard.

5.3.4 Roles

The satellite link can play different roles in the network:

- **Last mile connections**: End users are directly connected to the satellite to provide direct forward and return links. Traffic sources connect to the satellite feeder or hub stations through the Internet, tunneling, dial-up links, etc. The connection is basically from/to a gateway to a terminal.
  - For multicast, this is when the destination hosts are on the satellite network.
- **First mile connections**: The satellite provides forward and return link connections directly to a large number of ISPs' or other service providers' gateways, which will deliver the IP packet onward to the end users. As with the last mile connections, traffic sources connect to the satellite feeder or hub stations through the Internet, tunneling, dial-up links, etc.
  - For multicast applications, this is when the content provider host is on the satellite network.
- **Transit connections**: The satellite provides connections between Internet gateways or ISPs' gateways. The traffic is routed through the satellite links according to specified routing protocols and defined link metrics in the networks so as to minimize connection costs and to meet required QoS constraints for the given traffic sources.
  - In multicast context it can be combined with edge-casting and caching.

In addition to the above, a satellite network may serve a closed user group, and there may be intra-system connection, i.e. terminal to terminal.

Intra-domain connections may be open standards or proprietary solutions.

Inter-domain connections must be standardized and comply with equipment from multiple vendors.

It is worth repeating the BSM roles from TR 101 984 [4], as it is also valid for multicast.

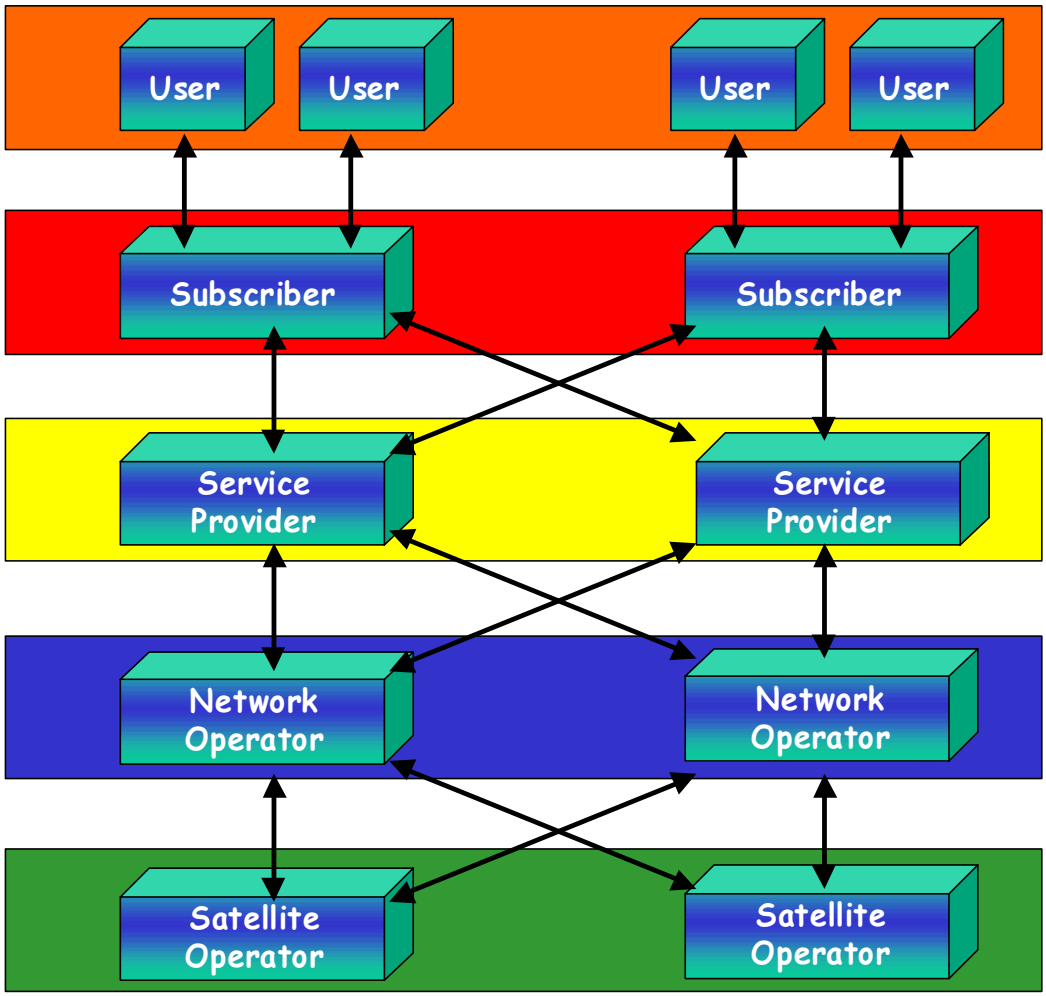

Figure 9: BSM roles

5.4 Session management

The text in this clause is based on IETF RFC 3170.

Session management is one of the most misunderstood services with respect to multicast. Most application developers assume that multicast will provide services like security, encryption, reliability, session advertisement, monitoring, billing, etc. In fact, multicast is simply a transport mechanism that provides end-to-end delivery. All of the other services are application-layer services that must be provided by each particular application. Furthermore, in most cases there are not defined standards for how these functions should be provided. The particular functions are dependent on the particular needs of the application, and no single method (or standard) can be made to be sufficient for all cases.

While there are no generic solutions that provide all session management functions, there are some protocols and common techniques that provide support for some of the functions.

With respect to session advertisement, there are a number of mechanisms for advertising sessions. One commonly used technique is to advertise sessions via the WWW. Users can join a group by clicking on URLs, and then having a response returned to the user that includes the group address and maybe information about group source(s). Another mechanism is the Session Description Protocol (SDP). It provides a format for representing information about sessions, but it does not provide the transport for dissemination of these session descriptions, nor does it provide address allocation and management. SDP only provides the syntax for describing session attributes.

SDP session descriptions may be conveyed publicly or privately by means of any number of transports including web (HTTP) and MIME encoded email. The Session Announcement Protocol (SAP) is the de facto standard transport and many multicast-enabled applications currently use it. SAP limits distribution via multicast scoping, but the current protocol definition has scaling issues that need to be addressed. Specifically, the initialization latency for a session directory can be quite long, and it increases in proportion to the number of session announcements. This is to an extent a multicast infrastructure issue; however, as this level of protocol detail should be transparent to applications. The session management service needs to:

- Advertise scheduled sessions;
- Provide a query mechanism for retrieving information about session schedules.

5.5 Caching

Caching is a technique, which can dramatically improve performance, whilst simultaneously reducing the traffic load on the network. A cache is a place where temporary copies of objects are kept. Once data has been cached, subsequent requests for it will be given the cached copy. Some Web browsers also implement their own caches on disk and/or in memory.

In a satellite environment caching can be done locally, in the satellite terminal, or more centrally, in gateways or hubs. The choice would depend on the application and the available technology. A larger unit serving more users, like a gateway, would normally have a larger set of resources available, like disk-space, and central caching would also not require satellite resources to access content to be cached, like for updating the content. However, every time a user accesses the cached content, the satellite segment would be used, and thus it might be a more efficient use of resources to multicast (or even broadcast) some data to all terminals capable of caching content if it was foreseen the content would be accessed by a large number of terminals (i.e. users).

There may be a combination of gateway and terminal caching, balancing the type of content.

With respect to terminal caching there may be some legal issues involved, like who has the right to use the storage on a device in your private satellite terminal, and is can content you have not asked for, including potentially offensive content, be stored there without specific permission.

Another issue is if technology should allow for caching multicast streams intended for real-time viewing (or display) for later use, or even for forwarding to others.

In a multicast scenario it may also make sense to cache multicast transmissions somewhere in the satellite network in case they need to be resent, for instance due to congestion or transmission errors.

Caching data offers benefits such as:

- Faster performance on cached data.
- Less network traffic generated.
- Less demand on servers.

There are also some drawbacks:

- There can be slower performance if the object is not cached.
- Data can sometimes get distorted if the target server becomes unreachable.
- Caching can confuse logging and access control.
- Not everything can be cached, such as a dynamically generated object.
- The cache server logs can be used to determine individual users' activity, possibly intruding on users' privacy.

However, multicasting over satellite combined with caching and the right management is probably a very strong and potentially beneficial capability of satellite systems.

5.6 Edge- and data-casting

Edge casting is a particular application of multicasting, where data is multicast and stored at the destination location. We would normally assume that there is a storage unit associated with the destination location, where the content can be retrieved at a later time. However, edge-casting may also be done "live", without any storage.

Edge-casting has several advantages over television broadcasting, as content providers are able to use e.g. Web technologies to track usage for billing purposes and to obtain information on customer behaviour. End users receive additional control over the viewing experience, including record and playback capabilities; on-screen, supplemental text content; and the ability to select a single camera view of pay-per-view event.

Satellites exhibit particularly strong capabilities to do edge casting. Edge casting over satellite can be split into two concepts:

- Caching in the terminal or terminal attached computers.
- Edge casting for local caching at gateways.

5.6.1 Edge casting to gateways

Edge-casting enables data like live video and audio to be up-linked to a satellite network in order to bypass a congested Internet backbone. Edge casting may be used, for instance, to distribute data to a number of gateways, or local access points, where it can be stored for local access. For instance, there is a product that caches the most popular television clips, like the latest news and the whether, at a local server, saving the cost and time of accessing these over the Internet. Likewise, the most popular Internet pages may be stored in a local gateway cache. These gateways are at the edge of the core network, sometimes called the edge network. All such stored data can be updated via multicast to the edge of the core network.

Figure 10 illustrates edge-casting to gateways, where one gateway distributes the content to the other gateways via satellite.

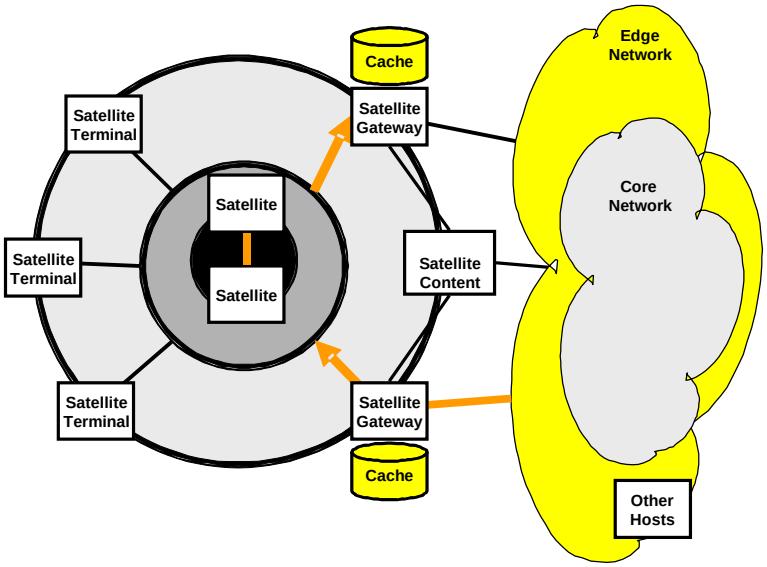

NOTE: This link can be via satellite or terrestrial.

Figure 10: Edge-casting from one gateway to another

5.6.2 Edge- and data-casting to terminals

Edge casting to terminals builds upon the same philosophy as edge casting to gateways. Edge casting to terminals could also imply transmitting corporate data to the edge of the corporate network. For instance would possible an auto-manufacturer be interested in edge casting the latest parts- and service manual for their autos whenever there was a change. However, storage capabilities would normally be lower at the terminal side than at the gateway side, and without proper filtering the majority of the content may perhaps never be accessed before it is outdated.

Data casting is a word put together from data and broadcasting, and as such it is equal to the broadcasting of data. Receivers can be tuned into specific types of data and always store the most recent version of, say, a web news papers or stock quotes. If broadcasting is used as a carrier for the data, one would not anticipate any return channels or acknowledgement of received data. All would be based on best effort. However, if the data is broadcast frequently, having the second most recent copy will in many cases not be critical.

Figure 11 illustrates edge-casting and data-casting to terminals, where a gateway broadcasts or multicasts data to a group of terminals, which in turn stores the content, possibly after some filtering, in a local cache. Caching the content enables a possible interaction to take place when and if the user accesses the content, which could be of significant interest e.g. for billing purposes.

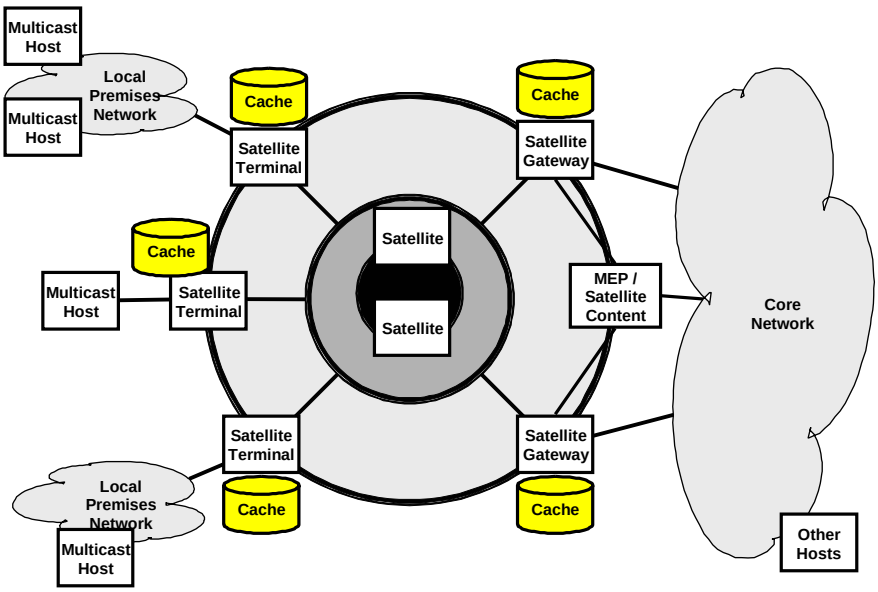

Figure 11: Edge- and Data-casting to terminals, which can cache the content

6 Satellite multicasting issues

6.1 Introduction

Different satellite systems exhibit a number of relative differences, which may present a challenge for harmonizing multicast standards. Such differences should preferably be taken into account for future standards, as they should allow for differences in technology. Some satellite system differences relevant for multicast may be:

- the total capacity they offer, both individual and total for the system;
- the market they target, business or consumer, and the global or regional span;
- the technology they base services on, including protocols, return channels and (possible OBP) satellite technology;

- the cost of the user equipment, service provider equipment, and the terrestrial network structure;
- the cost for the transfer of data, and also the billing models (flat rate or volume based);
- the frequency bands they operate on and thus also the fading characteristics and thereby probably also the error rates and the error occurrence patterns;
- the management concepts and the control plane.

The scope of BSM networks is defined in terms of four basic domains: the user domain, the access domain, the core network domain and the content domain as illustrated in figure12. The content would in this case specifically imply multicast content, i.e. a multicast source. A better word for the access network in this case would be the distribution network, as the traffic flow would be down to the satellite terminals. Please refer to TR 101 984 [4] for further details on the BSM architecture.

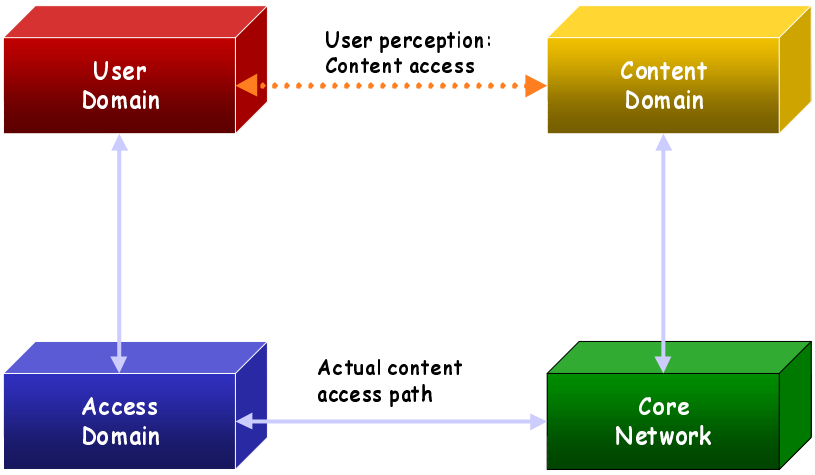

Figure 12: The four BSM domains

The four domains can be expanded into a more detailed diagram as given below. With respect to the multicast focus it is worth noting that the access paths may also be unidirectional, as will consider satellite forward channel with either:

1) No return.
2) Terrestrial return or return via another system.
3) Satellite return through same satellite network as in forward link.

In the present document, the focus will be the group of receivers that belong to one and the same satellite system.

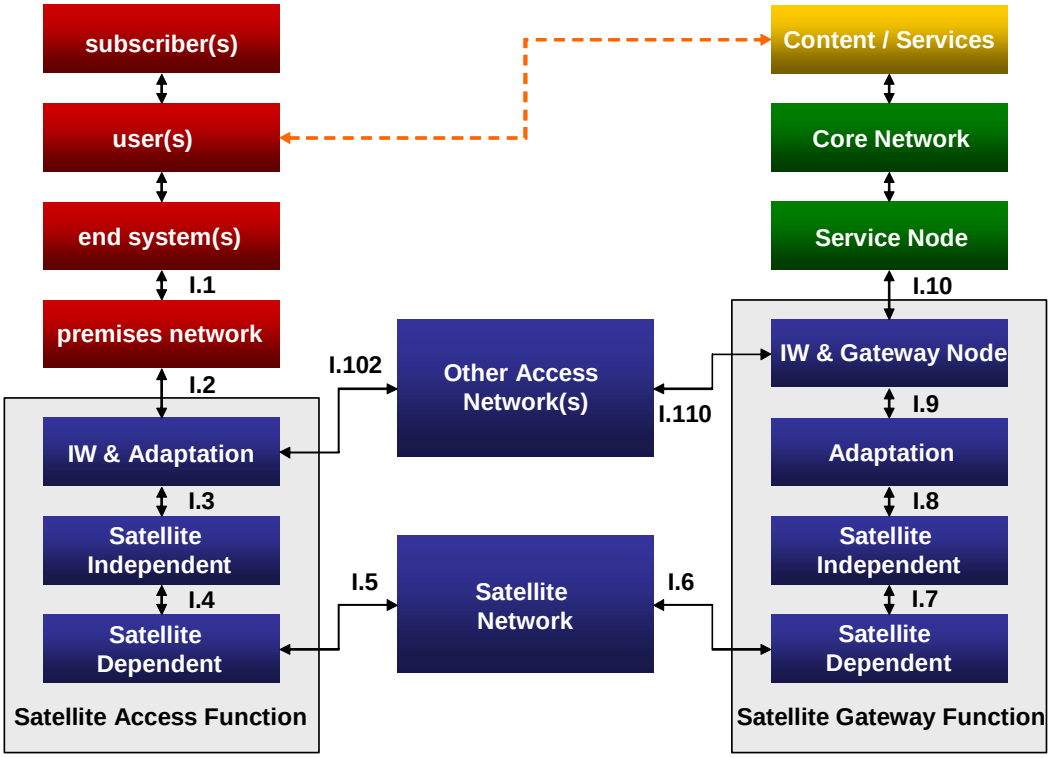

Figure 13: Detailed BSM reference architecture

The above reference figures are good for showing the different layers and functions in a unicast satellite system. However, we need to further develop these into a multicast network diagram. For this purpose we need to take into account the fact that there will be a number of terminals. Indicating three, as a minimum, can illustrate the fact that more than one but not all terminals receive data.

We also need to show that there may be more than one uplink point to the satellite network, and also more than one network entry point, e.g. more than one gateway. This is required to indicate the problem of selecting both where to uplink an incoming multicast source to the satellite network, and also that different gateways may be required to reach different terminals. More than one gateway is also required in a multi-satellite system, and more than one gateway may also be present for a number of other non-technical reasons.

Further, there may be more than satellite, and we have chosen to illustrate this by showing two. There will also be a network control center. For simplicity we also assume this will be the control center for the multicast network.

We need to show the networks that are or can be attached to the different satellite access nodes (gateways and terminals).

Figure 14 illustrates a simple satellite multicast concept with more than one gateway, more than one satellite and more than two terminals. Multicast flow is unidirectional, and in figure 14 from the right to the left. The terminals are simplified by not showing the interior at this point.

On one terminal there is more than one host. One gateway does not reach all terminals, nor does necessarily one satellite. More than one gateway or satellite may be in use for geographical reasons on a global environment, or for capacity reasons.

To summarize, figure 14 thus offers a generic satellite multicast architecture, and it seeks to illustrate the following aspects:

- There is a core network where multicast messages can arrive from.
- There can be more than one gateway. The network control center is associated with one of the gateways.

- There can be more than one satellite. In this case the satellite is to be understood as the satellite or space segment domain.
- There are a number of terminals, and these can in principle also originate multicast messages.

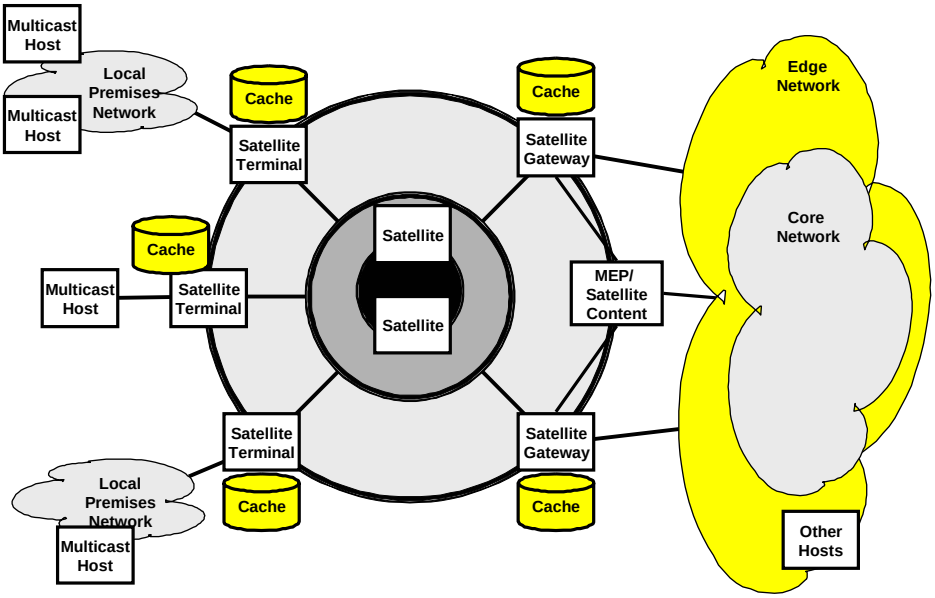

Figure 14: Basic satellite multicast concept showing different entities involved

Figure 14 does not explicitly show multicast originated within the satellite network, which of course can be the case. Internal sources will have to be considered in the following way:

- When satellites cannot replicate the message the multicast source must be tunneled to one or more gateways, and then we are back at the picture above except the source is not external. The routing problem is exactly the same.
- When there is more than one satellite involved, as in a global coverage system, then figure 14 will be valid for the satellite systems where the source does not originate.
- When the satellite does replication, then we simply look at the source as the gateway in figure 14. Terminology tends often to be a little different for bent pipe and OBP systems, and both the terminals and the gateways are in fact satellite access nodes. Thus the picture above is still valid, and other terminals can be reached with a lowest possible delay.

The simple picture above does not consider advanced inter-satellite routing mechanisms that for instance would be required in a LEO system directly, but assumes that once the multicast data is sent to the space segment, it finds the terminals that can be addressed via this or these satellite(s).

Given the number of technology options for the elements within the circles in the diagram it is the most important issue to define the interfaces for the boundary circle. The middle circle is the interface(s) between space and group while the inner circle is the interface between different space elements (satellites). The latter could for instance be inter satellite links or via gateways on the ground.

The following short-form terminology is used for satellite specific multicast issues:

- Satellite
  - The space segment, which in general can consist of more than one satellite. It includes the satellite control center. The space segment also includes the infrastructure required to operate it, such as the Satellite Operations Center (SOC).

- Terminal/also called Satellite Access function (TR 101 984 [4])
  - An earth station that does not provide terrestrial access to the Global Information Infrastructure for other satellite terminals. Terminals can be connected to internal "campus" networks, and can be both the source and the destination for multicast data.
  - In general there will be a (small) network connected to a terminal that in turn serves more than one user, and where every user in general can have more than one session.
- Gateway
  - An earth station that provides a gateway path for a number of terminals into the Global Information Infrastructure and handles traffic to and from multiple terminals. A gateway is normally operated by a network operator.

Note that there need not be any difference in the transmission capacity of a terminal and a gateway, although one may assume that in the general case a gateway will serve more users than a terminal. The air interface may also be the same for terminals and gateways if a regenerative satellite system is used, but for bent-pipe systems they are complementary in the sense that the gateway receives what terminals transmit and vice versa.

In some BSM systems with OBP, all satellite access node air interfaces may in principle be equal, although differences in overall uplink capacity (in particular) is expected. However, it is generally possible (and easy) with present systems to distinguish between terminals and gateways. With respect to multicasting, the major issue is whether the equipment has knowledge of the location of terminals that belongs to it (at the moment), whether is can or shall replicate messages and handle the lower layers of the link management. Many of these functions may reside in a network control center.

It is possible to build a BSM network as a L2 multicast-capable LAN or as a routed L3 IP network. If it is the latter it acts as a multicast router. Multicasting can be divided into primary functions, like:

1) the content source provisioning;
2) conditional access, including authentication, and key distribution; and
3) joining/pruning, and satellite spot beam and region beam replication management.

All of these functions need not be present in a given satellite system, or if they are, they need not be performed by a single terminal. But in the general case of a regenerative satellite with hundreds of spot beams and mesh connectivity the above are distinct and separable functions. The authentication and key distribution center may be located at a service provider gateway.

A group of receivers for a common multicast message is called a *multicast group*. A multicast group may in general include hosts everywhere in the world, and on a number of different networks. We therefore introduce the term *satellite multicast group* for the group of hosts that are part of a given satellite network. Every satellite host must therefore communicate (receive or send data) via a *satellite terminal*.

A satellite multicast group may be a subset of a larger group or a local group on the satellite network in focus. In any case the optimization problem is the same: How to most efficiently distribute the multicast data. Efficiency is here considered to be with the minimum use of spectrum resources while fulfilling the multicast requirements that are given for that group (i.e. time delay, etc).

In this report the focus is on the group of receivers that belong to one and the same satellite system, but where they are only a subset of the total multicast group in the Internet. In this way the satellite network needs to support multicast in the public internet, both sending and receiving, but can also provide closed "value added" multicast services specifically over the satellite network.

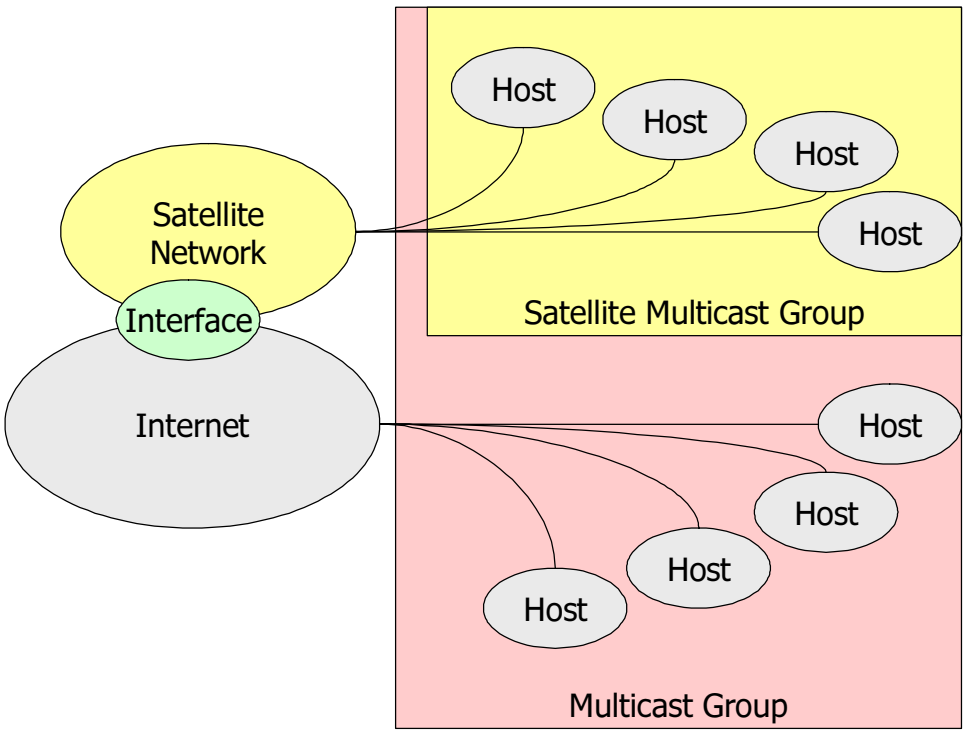

Figure 15: Satellite multicast illustration, indicating that the satellite multicast group is a subset of the total multicast group for the given source/content

The development and investment in broadband communications and networks over satellite in recent years has been mainly based on three approaches: bent-pipe, ATM or ATM-like fast packet technology, and DVB for broadcasting. None of these were originally designed to support IP multicast, but they have now been adapted to support IP multicast over satellites.

As networks evolve towards an all-IP solution, further options may need to be investigated at a research level: an all-IP satellite with on-board router. Such an option will need a significant amount of new system design, such as replacing the ATM and DVB switching with an on-board router, and will need to convince industry to develop and deploy satellite payload systems based on the new router technology instead of existing technologies. This may lose the benefits of ATM and DVB-S, which are already available. The benefit of an IP-router-in-the-sky approach is that the routing algorithm can be used to integrate the satellite links in an IP multicast routing tree at the source, trunk or end branch, as first mile connections, transit connections or last mile connections.

6.2 Multicast via Broadcast

Broadcast implies sending the data to every user in the network. Thus if there is a single subscriber somewhere on the satellite network the multicast message would be broadcast to all receivers in the network. The individual receivers must then sort out if the message is intended for them.

Multicast via broadcast implies a non-critical filtering of multicast messages in the gateway or satellite segment. The filtering must be done in the terminal segment.

Multicast via broadcast can for instance be done via an encryption key distribution, so that when a terminal (i.e. a host) subscribes to a specific group, it receives the present key for that group from a gateway/NCC, hosting the nearest network router. A broadcast channel has (or uses) no return channel.

Multicast via broadcast would be a very simple concept from a networking point of view, with a relatively static configuration. However, the concept is likely to pay for the simplicity with increasing the amount of forward channel resources and added processing complexity at the destination receivers.

Full broadcast does not take advantage of specific multicast benefits, will not be discussed any further here. However, in many BSM systems, depending upon how terminals are addressed, then the last part of the data delivery can be on a shared channel and the multicast data itself could be physically transmitted on a broadcast channel. The logical distinction is whether it is addressed to specific receivers or not.

6.3 BSM multicast capability issues

With respect to multicasting there are a number of application-issues that is worth considering when designing a satellite multicast system. Some of the most relevant are discussed below:

- Reliability, implying if every host needs to receive the data and if that information needs to be fed back to the network.
  - Usually reliability is required, but for some types of multicast applications over satellite it is not required to *know* that everyone receives the content. In other situation where using the latest data is crucial, the application may need to know that everyone is using the latest data. Examples of such may be stock-quotes.
  - If the information about the delivery needs to be fed back to the sender, then a return channel is required. A return channel can also be used for the implementation of an ARQ protocol that enhances reliability. However, if there is no return channel then error correction coding may be applied to reduce the probability of errors caused by the channel. Coding will however not prevent packets from being lost for instance within the operating system.
  - In some cases the application may need to be totally reliable, like in banking or some business or military applications. This may also have to do with QoS and the contracts with the users. There are however "soft-fail" issues, like if users subscribe to news or SW upgrades, where not having the latest version need not be critical, as it can be fixed later if required. If the content is multicast often, it will also probably be correct the next time, thus at any given time the probability of having the wrong content is small. In a streaming video application, for instance, an error in a video frame or in a few frames may be annoying, but not crucial. This should be configurable. It may be desirable to be able to adapt coding and power based on some criteria related to the reception of data. A satellite multicast standard should allow for a service provider to customize services, and to some degree also let users chose, but probably should not be too many different options/versions of the SW in any given network.
  - A BSM network must allow operators and service providers to set different levels of reliability, and should also allow the users in terms of QoS contracts request different levels of reliability.
- Timing and synchronization issues, including low delay delivery and the ability to constrain time differences between when hosts receive the multicast content. Also time-bounded delivery is relevant.
  - Synchronization of content reception time is usually considered important when users would have to respond immediately, like in gaming, where a rapid reaction could be required. For many game shows, it would be a requirement that data were revealed exactly at the same time at receivers. Absolute delays may be tolerable, but relative delays may not. In other cases it may be vice versa. For interactive applications, like multi-party video conferencing, the requirements should reflect that too large time differences could be annoying.
  - Low delays would be lower bounded by the satellite transmission delay. Over a one-hop satellite system, the relative variations in delay to different receivers need only be minor over the space segment. Over multiple (hop) satellite systems or LEO networks the delay may increase of several earth-space hopes are required, or even via inter-satellite links. Satellite networks are obviously not able to beat the laws of physics and offer shorter delays than the actual transmission time. However, what satellites can do in many cases is to offer a *predicable* and perhaps even a fixed delay, which may not often be the case in the best-effort ground network.
  - Time-bounded delivery is relevant when old application data is no longer valid. For instance, and old streaming voice frame is no longer of interest if a newer one has already been transmitted. If the data has this characteristic, the data could be checked before replication, and be discarded if no longer valid. To preserve spectrum resources, a satellite system should allow for such validity check, if required. More investigation is required to determine how this can be implemented, and to what extent it will be useful.

  - Satellite specifications should therefore allow for constraining differences in receive timing, a feature particularly relevant for multiple satellite hops. If the data needed to be synchronized to all receivers, it could perhaps be supported by taking advantage of system clocks present in the system. BSM satellite multicast specification should be able to specify/predict delay characteristics.

- Scalability and the ability to smoothly scale to large numbers of receivers.

  - For satellite systems the benefit of using multicast would increase with the number of receivers, and since the geographic coverage region would be large, one might assume also a large variation in interest groups, and thus possibly a large number of groups. However, there would also be some common topics of interest, and since the number of receivers can be large, both large and small interest groups can be foreseen. Further, new satellite technology, like larger antennas and more power in addition to using higher frequencies will lead to the ability to cover more users and thus potentially larger groups.

- The ability to provide ordered data is inherent in most satellite systems.

  - This would not be a big problem over the satellite segment, as the satellite segment itself would hardly influence the order. In some protocols the packets are numbered, and can be rearranged if they are out of order. In many applications, the higher level protocols would trigger a retransmission if there was a problem.

  - BSM networks should secure that the satellite network does not alter the packet order.

- Many-to-many multicast with more than one source.

  - This is a relevant problem in multi-party video conferencing. (Over satellite it may not be practical with a bent pipe satellite system, since the gateway would have to relay the data, and the delays would be relatively large for this type of application). In a mesh scenario interactive multiparty videoconferencing it is possible, but complex (in a number of ways). In many cases the multiparty scenario could also be modelled by a number of single-sender multicast scenarios. Efficient resource management could be a challenge depending on the nature of the data-streams.

  - Many interacting senders should be supported by BSM systems in the long term, but in the short term it would not be a first priority due to the complexity of handling it.

- Data flow intermittent applications.

  - Such applications could be the case, but most applications would probably be streaming type in the short term. For satellite applications the issue of radio resource management is important in this context, and this is as of today often solved individually for each system, which makes it challenging to define a specific way to handle intermittent flows in the short term. However, in the medium to long term BSM multicast specifications should support intermittent data flows.

- The BSM multicast solutions needs to work with applications in the public Internet with satellite systems both with and without a return path. It also needs the ability to provide secure data delivery.

- Caching/Storing capabilities should be included in the specifications, recognizing that cache sizes may vary significantly in end user equipment. In principle the cache could also be zero, which is relevant if a hard disk or other storage device is not present, full or broken.

6.4 Satellite multicast architectures

The basic reference figure 16 is applicable for different satellite network architectures. In consistency with TR 101 984 [4], satellite architectures are initially divided into mesh and star networks.

There may be a number of hosts connected to one and the same satellite receiver. This is an important observation, and one that can be exploited for satellite networks. In some environments, like a home environment, it might be desirable for all computers or TV sets (or peripheral equipment) connected to the satellite terminal to have access to the multicast content, without setting up a separate multicast address for each computer connected to the local network.

The general assumption is that the source is external seen from the satellite network. The special case of internal sources from satellite access terminals is discussed under mesh networks.

The general architecture core is shown below. This is the generic version which will be used to illustrate various aspects. The inner circle corresponds to the inter-satellite interface, the middle circle to the radio (air) interface and the out would correspond to the SI-SAP. The BSM network is all within the circles.

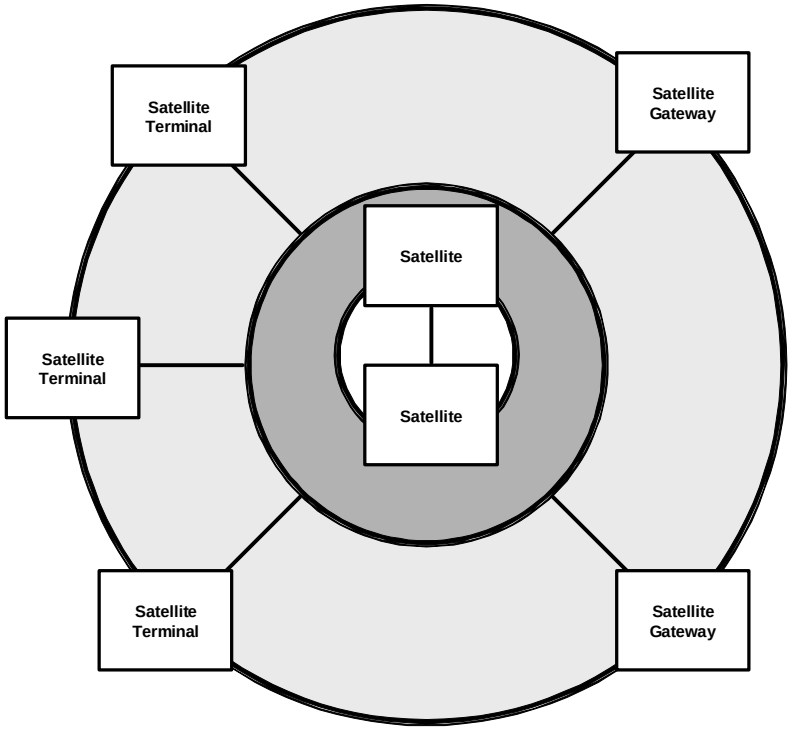

Figure 16: Generic satellite multicast architecture

A more detailed tree structure example is shown in figure 17, where the source is indicated at the top, and traffic flows down via gateways, satellites and terminals where switches or routers replicate the data according to the distribution tree. The purpose of this figure is to provide an illustration that perhaps may be more familiar in relation to multicast. However, it basically shows the same information as the simpler circular diagram above, and thus serves to verify the applicability of figure 16.

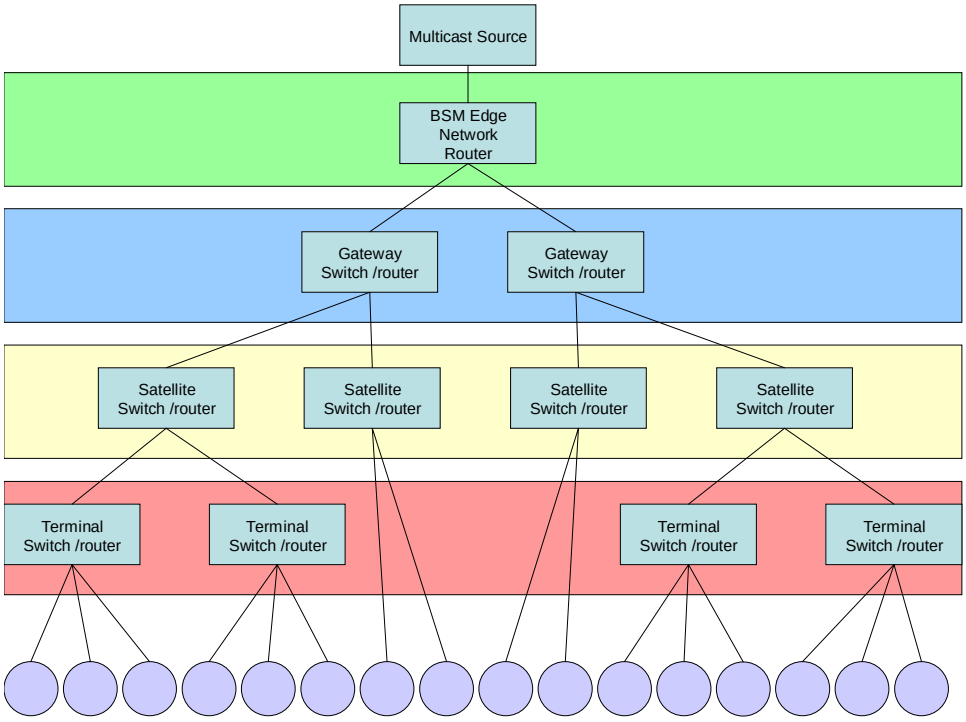

NOTE: Hosts are shown as circles at bottom and may be address directly via single user terminals or terminals with routers or switches. Different domains are indicated as horizontal colour bars.

Figure 17: Multicast distribution tree example

With multicast being a dataflow spreading out from a source to destination hosts in a BSM satellite system it appears the tree architecture is the most relevant representation from the source uplink and to the destination hosts. If several gateways are involved, the gateway routing architecture will depend on the interconnection, as the may or may not be connected via terrestrial means or satellite.

- Star satellite networks support simple tree based distribution from a gateway to satellite access terminals, in addition to star based routing.
- Mesh networks would also support a tree-based architecture, with branching in the satellite if replication takes place there.

A return interaction channel thus a no specific influence on the traffic distribution, but may for instance be used for group management and ARQ protocols.

As a comment on the Steiner Tree problem mentioned in clause 4 related to general multicast routing, we will claim here (without proof) that the general multicast routing problem limited to within a BSM network (with a GEO satellite) is significantly simpler, and in general not an NP-hard optimization problem, but at most a Shortest Path problem, and more often the well studied Minimum Spanning Tree. The specific mathematical problem will depend on the actual satellite network architecture. However, in a multi-gateway scenario and for external source the problem can again be a Steiner Tree problem.

6.4.1 Star networks

6.4.1.1 Single gateway/hub

For convenience we repeat the star network reference architecture from TR 101 984 [4], to show how it maps onto the multicast reference network architecture.

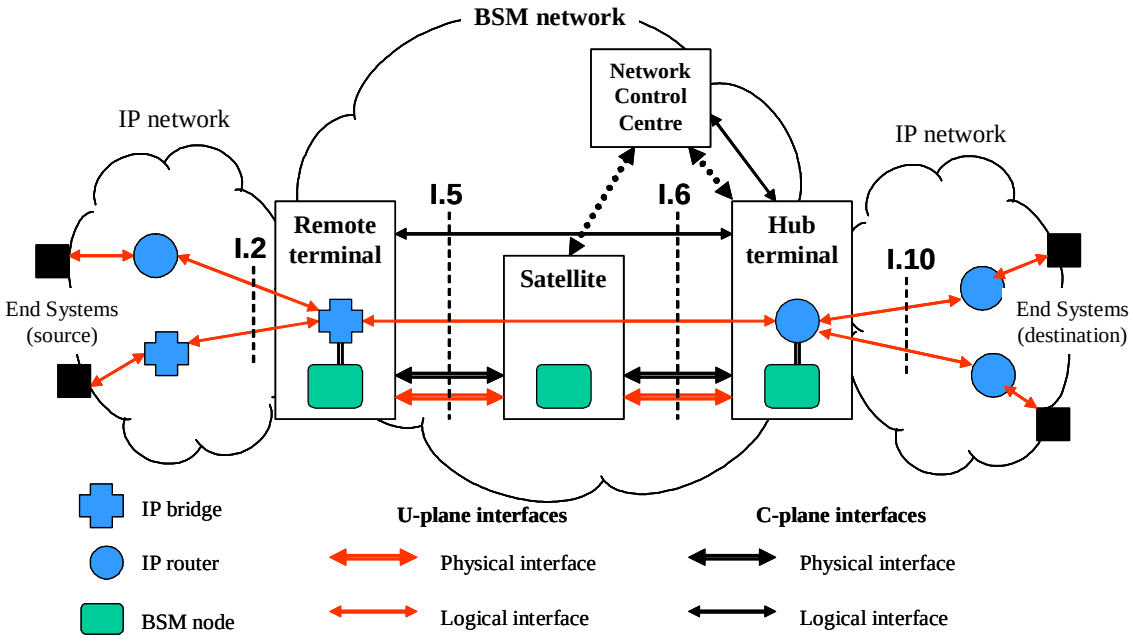

Figure 18: Star network architecture from TR 101 984

The purpose of figures 19 and 20 is to illustrate the routing principles with more than one addressed terminal and the aspect of the location relative to different beams for the end terminals. The generic figure has now got the traffic paths highlighted. It shows traffic flows from one gateway to all terminals with one red uplink. The figure to the left shows graphically the gateway-satellite-terminal situation, and indicates the simple routing tree that is valid in this case. All multicast terminals are reached by the same downlink (illustrated by all arrows having the same colour).

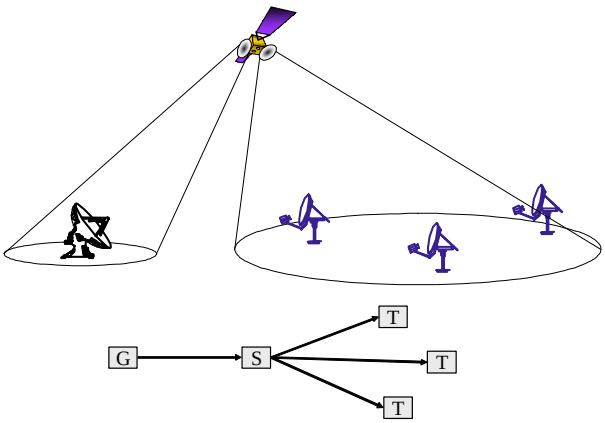

Figure 19: Star Architecture, single Hub, global beam as in a traditional bent pipe system for broadcasting

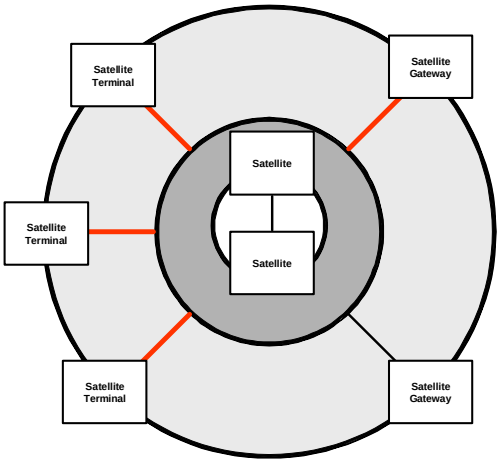

NOTE: This figure covers the bent pipe global beam and the multiple spot-beams with the replication in the satellite.

Figure 20: Star network mapped onto the multicast reference architecture

With a single hub as the contact point for the external core network, addressing from the external world is quite simple. The hub would be the Contact Point (CP) for all the terminals (and the hosts connected to them) for that satellite system. This architecture is found in DVB-RCS systems.

All multicast data to any host connected to any terminal on the satellite network would have to enter the satellite network through the CP node. This applies both for layer 2 and layer 3 systems. If there were no spot-beams, then the content would in turn be sent in the downlink global beam and the terminals that subscribed to that content could receive it. The ability of a given terminal to receive data in a beam also depends upon the number of carriers (i.e. transponders), the address space and processing capability (e.g. number of PIDs in DVB).

For example, if there are spot-beams, then the hub may or may not need to send the content several times, depending on the satellite architecture, and if the satellite is able to replicate data. In the case below a bent-pipe satellite requires two (or more in the general case) uplinks to reach all destination hosts in two spot-beams. Thus a copy of the multicast content must be present at all gateways.

In any case, the single hub architecture can be illustrated with figures 21 and 22. This figure illustrates one gateway, as the sole contact point for the network (the other gateway in the figure has no transmissions from it in the drawing), a satellite segment, which again could contain one ore more satellites, and a set of satellite receivers. In figures 21 and 22, two copies of the uplink multicast data is required, as the satellite does not replicate the messages, to reach the (three) multicast receiver terminals. Two terminals are shown in the same red beam.

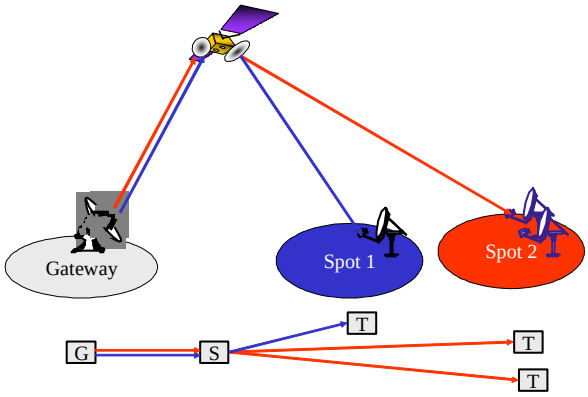

Figure 21: Single gateway, multiple spot beams

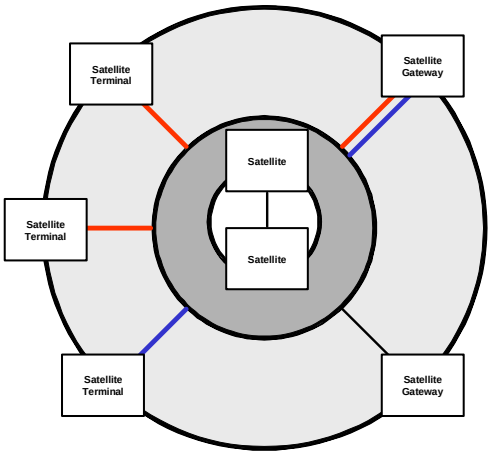

NOTE: There are no uplinks from the lower gateway, so it is void.

Figure 22: Single gateway, multiple spot beams

Acknowledgements (in a reliable multicast scheme) or replies to IGMP messages could potentially congest the gateway receiver if these were sent from many locations simultaneously, and a mechanism that resolved these issues would be required. Such a mechanism could for instance spread the return messages out in time.

For the star network, the hub would have to:

- Maintain a set off tables for which satellite terminals need to receive which multicast group.
  - Optionally distribute new keys when a member joins or leaves a group.
- Track acknowledgements from hosts.
  - Optionally resend, when required and depending on protocols, content that was not received well.

If there were multiple spot-beams, yet all accessible through one and the same gateway, that gateway would have to determine which beams should downlink the content and which hosts were connected to those beams, so as to save on spectrum resources. For the star network, with multiple spot beams, the hub would have to maintain a set off tables for:

- Which hosts subscribe to which multicast (source, group). (One table per group).
  - Maintain a table of which beams and receivers hosts belong to. Could be dynamic.
  - (Re-)distribute new keys when a member leaves a group.
  - Distribute keys when a member joins a group.
- Track acknowledgements from hosts.
  - Maintain a table of which coding/modulation/power should be used in each beam.
  - Resend, when required, content that was not received well.

6.4.1.2 Multiple gateways

When there are multiple gateways, an immediate issue becomes the choice of which gateway should be used to transfer the multicast content to the end receiver (and host). This consideration applies both when more than one gateway can reach a given receiver and when only one of the gateways can. Choices of which gateway to use can be either static or dynamic. The gateway choice is in general relevant both for external multicast sources and for internal multicast sources. For internal sources that are to be received in terrestrial subnetworks, gateways are used to forward the content to other, external (PIM-SM) routers.

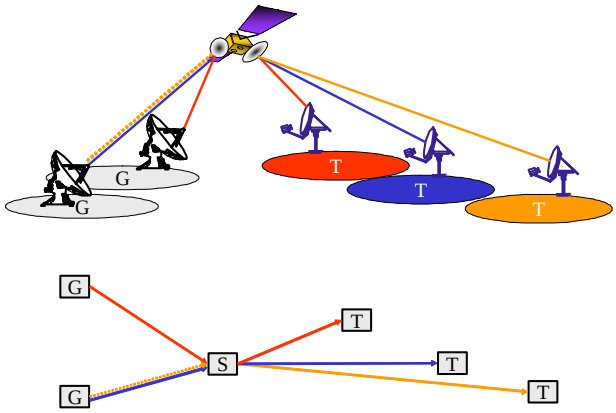

NOTE: This is a commonly expected scenario. The dashed orange line is present if the satellite does not replicate messages.

Figure 23: Multiple gateways and multiple downlink spot-beams

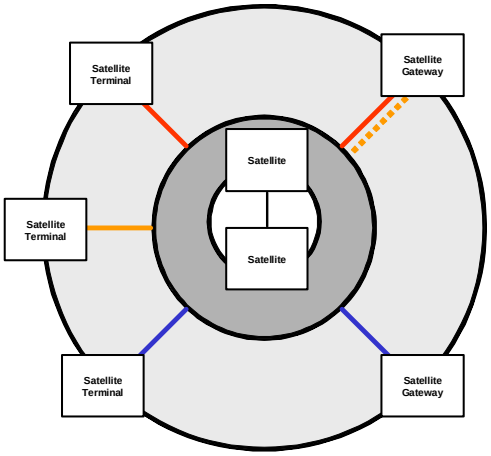

Figure 24: Multiple gateways and multiple downlink spot-beams mapped onto the circular diagram

The gateways could contain a router, and one of the gateways or another network node could be chosen as the logical **Multicast Entry Point** (MEP) into the satellite network. The MEP would be the satellite system multicast server. There would normally be a Network Control Center (NCC), and the NCC could, but need not, act as the logical MEP. The MEP need only be the first contact point for addressing the destination receiver, and it need not be a single entry point for all traffic. In fact it would normally not be the sole traffic entry point. Edge routers would be found at these entry points.

The MEP may be on a geographically distant location from where the traffic actually enters and/or exits the satellite network, as the satellite system might be global, but with one MEP.

There may be multiple gateways even with an (active and) regenerative satellite. This could for instance be due to uplink capacity reasons, or geographical (or political) reasons. Figures 23 and 24 illustrate the concept of spot-beams and multiple gateways. The choice of gateway for a specific host may be permanent or temporary. If not permanent it may be due to the host (e.g. a laptop) being moved, or because there is a dynamic resource mapping between the gateways and the terminals depending upon for example traffic loads in the system. In any case, with multiple gateways in the system, a choice of gateway must at one point be made, and the NCC could be the location that kept the host/gateway mapping.

In a multiple gateway scenario, the system needs to:

- Maintain a host/gateway mapping. This could be static or dynamic.

The system/gateway needs to:

- Maintain which hosts subscribe to which multicast group.
  - Maintain a table of which beams and receivers hosts belong to. Could be dynamic.
  - Optionally it may distribute new keys when a member leaves a group.
  - Distribute keys when a member joins a group.
- Track acknowledgements from hosts.
  - Maintain a table of which coding/modulation/power should be used in each beam.
  - Resend, when required, content that was not received well.

6.4.1.3 Internal source

If the source is internal then the content must be tunneled to a/the gateway. Once it reaches a gateway then the above considerations apply. There may be restrictions as to if the source shall be made available outside of the satellite system.

6.4.2 Mesh networks

Figure 25 shows the BSM mesh network figure from TR 101 984 [4]. The basic difference is the emphasis on inter-terminal communication without the need of going via a gateway, which is beneficial when a large portion of the networks traffic is between terminals on the same system, etc.

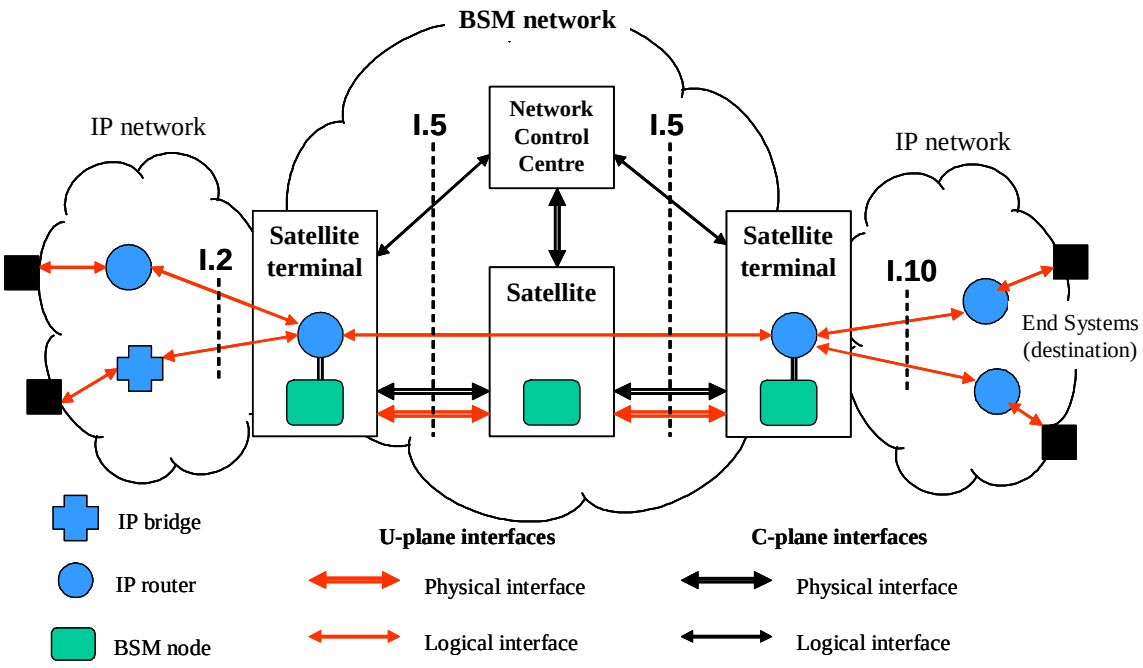

Figure 25: BSM mesh network reference

If figure 25 is developed into a simple mesh multicast figure, as below, we see the source indicated at the top and the capability to reach other hosts (circles) via other satellite terminals or external hosts on the public internet via the gateway(s).

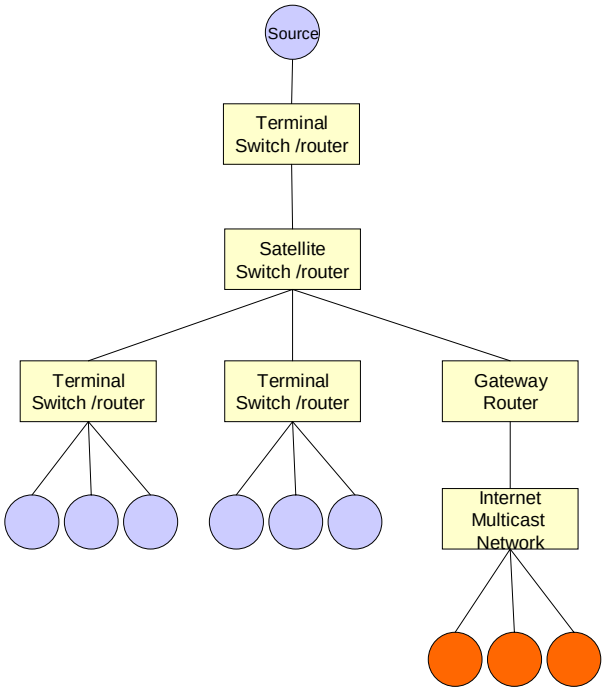

Figure 26: Simple mesh multicast network

6.4.2.1 Single gateway

The same terrestrial routing arguments as for the bent-pipe satellite is valid, except that in addition a regenerative satellite able to replicate data could do with one uplink even for multiple spot-beams, as the replication could be done in the satellite. In this case, one would save uplink capacity, as it would not be required to uplink the multicast data one per beam. An interesting advantage of satellite replication is that the gateway does not have to maintain all that transmission information about the receivers.

Reference figures are as for the transparent case for a circular diagram, with emphasis on the ability of the satellite to replicate messages (one uplink to reach all terminals).

6.4.2.2 Multiple gateways

Basically the same arguments are valid here as in the case for bent-pipe satellite systems. If there are several gateways then there must be a mechanism that chooses which gateway shall be used to reach which terminals. A possible additional issue would be that with a mesh network most gateways accessing the same satellite could probably reach all terminals in a technical sense. Therefore distribution trees could be more dynamic. However, in practice there may some issues that reduce the practical dynamics, as operators or service providers may always use their distinct gateways, for instance at a national level. There are also issues related to lawful interception.

Reference figures are as for the transparent case for a circular diagram, with emphasis on the ability of the satellite to replicate messages.

6.4.2.3 Internal source

A mesh network with the underlying OBP capability can in principle support multicast with one terminal as source and other terminals as recipients (for instance distance learning or corporate multicast). It is however worth noting that there are systems with OBP capabilities that cannot support direct multicasting, so the OBP is a requirement but not a sufficient requirement alone. Sufficient buffering and replication capabilities need to be present as well as a resource management system that handles the issues involved.

When satellites replicate multicast messages it is more likely in the short term they will do that a layer 2 than at layer 3, but both options are possible. It is however not so likely that a standard PIM-SM router will be implemented. A spoofed system tailored to the specific satellite in question capabilities is more likely.

A BSM satellite system may have other forward channels than the BSM satellite channel, like terrestrial options. A BSM system may also in principle be without a return channel, although the more likely scenario would probably be that if the return channel was not via (the same BSM) satellite, then it would be via a terrestrial channel, like ISDN.

6.4.3 Forward and return channels

An interactive satellite system will in principle have two channels, one forward and on return. In DVB-RCS the interaction channel is defined as a duplex channel in addition to the forward DVB channel, but in most cases the forward interaction channel will be physically combined with the forward (DVB TV) channel.

6.4.3.1 Forward channels

A *traditional* satellite system consists of one or more geo-stationary satellites with a bent-pipe payload operating in a global beam environment. In a one-hub scenario, all users can be accessed via the same uplink beam, although they may be tuned to separate downlink carriers.

Some satellite systems use multiple spot-beams on the downlink, and sometimes also on the uplink. In this configuration, users in different spotbeams *must* be accessed through separate downlink *carriers*. It is unavoidable to require a separate copy of the multicast data in every *spotbeam* where there is a receiver, but should if possible avoid more than one carrier in that beam. Efficient satellite multicast should however not send the data in beams where no target receivers are known.

The data may (or may not) have to be up-linked via different gateways to reach different downlink beams, depending upon the satellite architecture. Different spot-beams *can* be accessed via one and the same gateway, or hub, if there is one uplink beam or if the satellite has a switch on board. The uplink data may be either replicated in the satellite. If not, replication may take place in a gateway node. If a number of uplinks are required, the data may be multicast first to all gateways, as in a tree structure. Note that if multicast data is sent in all beams without considering if there are subscribing receivers there, the methodology is equal to broadcasting the data.

If the origin of the multicast is a satellite terminal (on the network in question), the same principles apply. A given terminal may or may not be able to reach all other terminals on the network directly, depending on the satellite architecture. The replication may therefore take place in the satellite, or the multicast data may be relayed via a gateway node. It is also possible, but perhaps not practical, that replication can take place in an originating terminal.

With respect to multicasting we can differentiate between the two scenarios:

1) All terminals can (and shall) be addressed via one specific contact point.
   - The point can be a gateway that replicates as required and reaches all terminals; or
   - The point can be a regenerative satellite that replicates and which can reach all terminals.
2) All terminals cannot (or shall not) be addressed via one specific contact point.
   - This usually implies that the multicast data must be copied to a number of different gateways in a single satellite system.
   - In a multiple satellite system, the data may need to be copied to different satellites.
   - There need to be some network node that decides which gateway that shall be used to reach any given terminal (and host).

6.4.3.2 Return channels

Return channels can be used in multicast when reliable multicast is required, but they are not mandatory. Satellite multicast delivery may or may not require a return channel, and the return channel, if present, may or may not be via satellite. If it is via satellite, it may even be over a different satellite from that of the forward channel. However, in an interactive BSM system, one would generally assume a return channel via the same satellite system. Still, we can thus define the following scenarios:

1) No return channel.

2) Return via separate system, including terrestrial fixed and separate satellite.

   - The return channel is non real time or not directly involved in content delivery.

3) BSM return via [integrated] satellite system.

   - The return channel is a real-time interaction channel directly involved in the content delivery.

BSM multicast should be able to work in all cases, and a multicast manager will decide the required actions depending upon the situation. Obviously, return channels can play important roles in securing delivery, in interacting with group management, and so on. In the case with integrated BSM return, the interaction channel can also be used (with ARQ) for securing reliable delivery. To some extent this must be compensated for in the other situations with efficient coding or packet duplication. With a terrestrial return, retransmissions may optionally be requested at the application level.

In the second case the return channel is used for non-real-time services. In fact there will in principle always be a return channel in some form either it is via satellite, via terrestrial data, telephone or mail or post. At some point the end user needs to sign up (or off) for the multicast groups, and the content delivery paths need to be set up. To a large degree the return path here decides the dynamics and granularity of the control loop at which this management can happen.

Table 1 indicates some of the services the different options can support.

Table 1: Different multicast options that are supported via different return channels

Service/use	No return channel	Non-BSM return	BSM return channel
Group join/leave	Email, web, phone, mail	Email, web, phone, mail, IGMP possible	Email, web, phone, mail, IGMP, PIM-SM
Multicast source	No	Not via BSM	Yes
Multicast receive	Yes	Yes	Yes
Reliable multicast	No, not likely	Perhaps, but not easy	Yes
QoS	No	No, not likely. Perhaps at application level	Yes
Dynamic groups	No, not for very short delays	Yes, depending on the delay	Yes
Interactive multicast	No	No, not unless return delay short	Yes, possible, limited by path delays

It is proposed to group the use-cases of interaction channel related to multicast into the three cases mentioned above. *At the multicast level, it is not crucial if a return channel is via satellite or not,* as long as it meets certain requirements that relate to the capability to be used interactively for content delivery down to the link level.

Low speed return channels with long latency may be considered not present related to link management. They have either to long delay or too low capacity (or both) for active, interactive use. However, the signing in and out of multicast groups, for instance, need not require high speed or low latency. (It could even be managed by email, post or fax, as when one signs up for pay TV channels).. With a fairly low latency channel one can offer more dynamic group management. There is however a principle difference as to if the application itself has the capability to directly sign up for new services or not (as in the fax or email case). If not, group management needs an operator intervention.

The general BSM system will have an interaction channel, and thus BSM multicast will in general offer the capability to be fully interactive. Further, most BSM system will also have a high speed and low latency (limited by the GEO arc time delay) channel, albeit some may have (or make available for multicast) a too low capacity for heavy interactive use. In any case the return channel can often be considered as an administration channel, and thus its use should be minimized.

6.5 User traffic forwarding functions

This plane would compare to the U-plane. BSM multicast would base the traffic at the link and physical layers on whatever technology was used also for unicast. This domain would be in the satellite dependent region of the protocol stack.

In BSM multicast some of the issues involved would which protocols to use (UDP will in the short term likely be the protocols into and out of the BSM system), the use of spoofing or interception and how and where to perform replication. The ITU is for instance working on an ECTP (Enhanced Communications Transport Protocol) (see clause 7.3).

It would be beyond the scope of work to propose any new transport protocol for BSM systems, but it is important to recognize that a BSM multicast system could in the future need to interoperate with other transport protocols than UDP.

A use of ARQ would be relevant to secure reliable data transfer. Therefore this issue is also discussed.

6.5.1 Satellite multicast replication concepts

Every multicast message must be replicated to every relevant downlink carrier (which in general would be one per relevant spotbeam). If all downlink carriers are addressed via one and the same gateway in a bent-pipe satellite system, then replication and mapping onto relevant carriers can take place there. If the satellite is OBP and not transparent, the procedure will depend on its ability to replicate data. If the satellite replicates then the problem is moved there, and the terrestrial handling is simpler. However, if different spot beams are addressed via different gateways then the terrestrial network must route data to the relevant gateways. Further, if different subscribers are handled from different gateways (independent of the technical capabilities, but based on i.e. subscriptions) then one risks sending the same message twice in the same beam (on even on the same carrier) unless the situation is handled.

A key element in a satellite multicast system is the replication node, or the replication nodes. This is where the message is copied once for every path to the next multicast router, or if it is the end router, where the message is copied one to every receiver host. This is also the key element that is able to take advantage of the satellites capability to reach a large number of receivers. Therefore a replication node must be able to efficiently address as many receivers (of multicast content) as possible with as few copies of the data as possible.

For a satellite system this implies that all receiving hosts in a downlink beam should (preferably) be addressed at the same downlink carrier at the same time. Some beams may have more than one carrier. In this case it would be a challenge to group the receiver hosts on the same downlink carrier. Note that if there are only a handful of receiver hosts it is of significant importance. If however there are a large number of receivers on each of the (possible) carriers then there is already a large gain obtained on both (all) carriers. In this case it may also be more difficult to move receivers to another carrier, but this would depend on the Radio Resource Management (RRM) system.

If the satellite has the capability to replicate messages the system can obtain maximum efficiency in terms of spectrum use, as it may in this case be possible to reach all receiver hosts with only one uplink copy of data even in a multiple sport-beam system. Further, the uplink could come from anywhere, specifically even a terminal.

On the other hand if the system was a bent-pipe system, a separate uplink copy of the data would normally be required for every downlink beam. More than one copy could in this case be required form the same location (gateway).

In most systems (today) the downlink to terminal is a TDM, and with the identified scenario where multicast content commonly would come from the Internet, then the TDM would originate in a gateway. A gateway may uplink more than one TDM to a spot beam, depending upon the bandwidth of the carrier and the available bandwidth in the spot-beam. This gateway would than be challenged to group all receivers to the same downlink TDM and command all receivers to read the same timeslot.

This challenge is in principle straight forward, but certain issues need to be considered, like:

- How is group management handled when some receivers leave a group? How are keys distributed and redistributed. How frequently can such groups me updated with careful use of resources.
- How is return traffic managed, if it is present? When shall which terminals respond?
- Shall there be any ARQ protocol?
- What shall the underlying protocol be (ATM, DVB), and what is the transport stream and segmentation?
- Is the transport connection oriented or connectionless?
- How is intra-domain routing handled compared to inter-domain routing?
- Where are messages that originate within the network with destination hosts also within the same satellite network replicated?
- How are real-time and timing constraints requirements handled?

The list above is not meant to be complete, but to give an indication of some of the issues involved. This is in contrast to broadcast as of today over satellite where nothing is optimized for or addressed to any particular receiver.

As always on radio and satellite transmission systems, the spectrum resources need to be used wisely, and not wastefully. Satellite multicast technologies should minimize the amount of administrative traffic, the number of back-and-forth signalling paths, and use a protocols that are which is robust with respect to delays and transmission errors. As a consequence satellite multicast should address all target receivers in the same beam through the same downlink message. When a multicast message enters the satellite system it should be replicated once to every beam where there is one or more destination hosts. Replication for distribution over the satellite link can be performed several places, such as in the satellite itself, in a network control center, or in one ore more gateways. If the message needs to be replicated in the destination terminal for the purpose of being received by several local hosts, this can be done by a local router. Thus, the replication can occur at a) the source, b) a gateway performing a redistribution function or c) the satellite.

6.5.1.1 Replication externally

External replication will involve a number of messages sent as unicast or point-to-point within the satellite system.

The replication can in this case be assumed to take place outside of the gateway(s) in the system. In this case a challenge for the satellite system would be on the resource optimization side, if at all possible. The system could risk receiving a large number of equal messages to different addresses. If no action is taken this would waste spectrum resources compared to minimizing the number of copies sent.

Whether the gateway actually would be able to recognize that there were multiple copies of the same message to different addresses but within the gateways same beam will depend upon the information the gateway has available. We cannot assume the gateway in general buffers messages and compares the content. Some assumptions have to be made, for instance that messages would have to arrive within a certain time interval and have the same content.

External replication is not considered as part of a satellite multicast solution.

6.5.1.2 Replication in a gateway

This is a very relevant scenario. Replication in the gateway would be a common scenario for satellite multicast, and in particular for a bent-pipe satellite.

- The gateway would need to maintain a host/beam mapping and replicate the message once for every beam where there is a receiver for the multicast message.
- The gateway would need to receive and act upon acknowledgement, if present.
- The gateway (or MEP or NCC) would need to maintain a list of hosts and multicast groups.

To optimize the resources, the gateway would have to group multicast receivers onto the same downlink carrier. In some beams there may be more than one downlink carrier. The terminals would also have to be told to listen to specific timeslots, if applicable. Thus, the gateway would need work across several layers, from the multicast application layer to the link layer and the link management and resource management plane.

The gateway may, depending on potential QoS settings for the multicast data, set specific coding and modulation parameters for the receivers. The gateway may also choose to implement specific acknowledgement schemes for reliable multicast.

6.5.1.2.1 Replication in a satellite terminal

A terminal may have a number of hosts connected to it, forming a home network, or a customer premises network. There is a chance that more than one of the connected hosts subscribe to the same multicast group, in particular that may be the case on a closed network (that may be part of a corporate network).

If more than one host connected top the same satellite terminal subscribes to the same multicast group, then the message should be sent to that terminal only once, and the satellite terminal should replicate it.

IP multicast receiver hosts could be connected to a satellite terminal using a multicast capable protocol (e.g. Ethernet) and a common interface. This could eliminate the need for replication in the satellite terminal since this is taken care of in the LAN.

If IP multicast receiver hosts connect to a terminal through several interfaces or the protocols in use are not multicast capable (e.g. PPP encapsulation), then the terminal should replicate the multicast packets and forward so that each local group member may receive an instance.

Acknowledgement of the received message, if required, can be either based on the satellite terminals successful reception of the message, or on the individual hosts' successful reception of the message. It seems to make most sense to let the satellite terminal resent, if required the message to hosts if required.

6.5.1.2.2 Replication in a multicast entry point

When the source is external to the BSM network, it has to enter the BSM network via one ore more Multicast Entry Points (MEP) somewhere. The same consideration applies for also points to point source, with the potential difference that it always would enter at one single place (the Network Entry Point, NEP). A multicast entry point could be the same as the network entry point. The NEP and the MEP does not have to be a fixed location. In a network with many gateways any one could serve as the source traffic entry point. MEP. However, there must be a node somewhere, somewhat similar to the MSC (Mobile Service Center) and HLR Home Location Register) in the GSM system, that knows which gateway to use to reach any set of multicast receiver terminals. A popular group could for instance have receivers served by most or all gateways in a BSM system. A less popular group or a regional interest group could have receivers served by only one gateway. The node that distributes traffic among the different gateways and the different satellite transponders is often called the Network Control Centre, NCC.

If we view the BSM network as an autonomous system, then an external facing border router (with an exterior gateway protocol like BGP) would be the NEP. The NEP could be co-located physically with the NCC (which is facing inward in the BSM network, and would i.e. run IGP). For multicast sessions a likely protocol set today would be Multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP). The MEP would host a multicast border gateway function and an internal and external routing policy.

Replication in the Multicast Entry Point (MEP) or Network Entry Pont (NEP) is relevant with respect to replicating to all the gateways in the BSM network, but it is also possible (basically for smaller groups, i.e. few subscribers) that the source is replicated to all downlink beams in the MEP, bypassing a need to replicate in the gateways.

The MEP could consist of (PIM-SM) router functionality, and replicate, as a minimum, the data to the different gateways that required them. Optionally it could replicate for terminals, but this would not be in line with a philosophy of replicating as late in the distribution tree as possible and not take advantage of the satellite multicast capabilities.

If the gateways that needed to forward the multicast data could communicate with the MEP via satellite then the MEP could use satellite to distribute (perhaps even multicast) copies of the message (adding path delay to the overall delay budget) to the gateways. Alternatively a terrestrial network would interconnect the gateways. A terrestrial network could e.g. be the public Internet or fixed/leased lines.

Considering multicasting in the forward link:

In an IP BSM network replication should be done as necessary to offer IP multicast packets to multicast group member terminals. Starting at the terminal this means that the IP multicast packets must be available in the forward link to the terminal. This leaves two options when a terminal requests traffic for a specific multicast group: Either provide an instance of the IP multicast group packet stream in the forward link to the terminal or switch the terminal to a forward link where there already is an instance of the IP multicast group packet stream.

Considering multicasting in the return link:

In an IP BSM network replication should be done as necessary to offer IP multicast packets to "multicast exit points" and multicast group members connected to satellite terminals. Looking at the "multicast exit point" this means that the IP multicast packets must be available at this point, provided through routing, switching and replication of packets provided by the terminal connected to the source. Looking at the terminal this means that the IP multicast packets must be available in the forward link to the terminal. This leaves two options when a terminal requests traffic for a specific multicast group: Either provide an instance of the IP multicast group packet stream in the forward link to the terminal or switch the terminal to a forward link where there already is an instance of the IP multicast group packet stream.

Where to replicate depends on the following factors:

- The topology of MEPs, "multicast exit points" and satellite hubs.
- The topology of satellite terminals with multicast group members connected, satellites and satellite hubs.

- The multicast capability of the protocols used to interconnect the units.
- The replication capability at the different units.
- The configuration capability of replication and routing/switching at the different units.

Generally, replication should be done close to the IP multicast group member, as this saves bandwidth.

6.5.1.3 Replication in satellite

For any satellite system, the optimal configuration in the spectrum sense is if the message can only be up-linked once, and then reach all destination receivers either directly or via replication in the satellite.

This can be the case for a global beam system with one gateway. For a spot-beam system, the most efficient use of resource occurs when the satellite itself is able to replicate the multicast message. Replication can be at different layers.

6.5.1.3.1 Replication at MAC layer

If the satellite replicates the message at the MAC layer it will basically copy the content of the up-linked frame to different beams. The satellite may in general need to buffer the content, if it cannot copy it instantly to other beams. This would depend on the traffic load and the RRM of the satellite, but in general buffering would be required.

Figure 27 illustrates an IP multicast distribution tree for the group G1. In this case, the satellite network is acting as a layer 2 bridge. The downstream satellite access terminals (not shown - part of the network cloud), with hosts and/or routers locally attached, are forwarding the multicast stream onto their local LAN connections (bold arrows). The SAT attached to the multicast source (in this figure via the router R1) is transmitting the G1 data stream into the Satellite network using the Satellite Multicast Group ID (MGID) allocated to group G1. The multicast distribution tree is constructed either statically or dynamically (using multicast routing protocols).

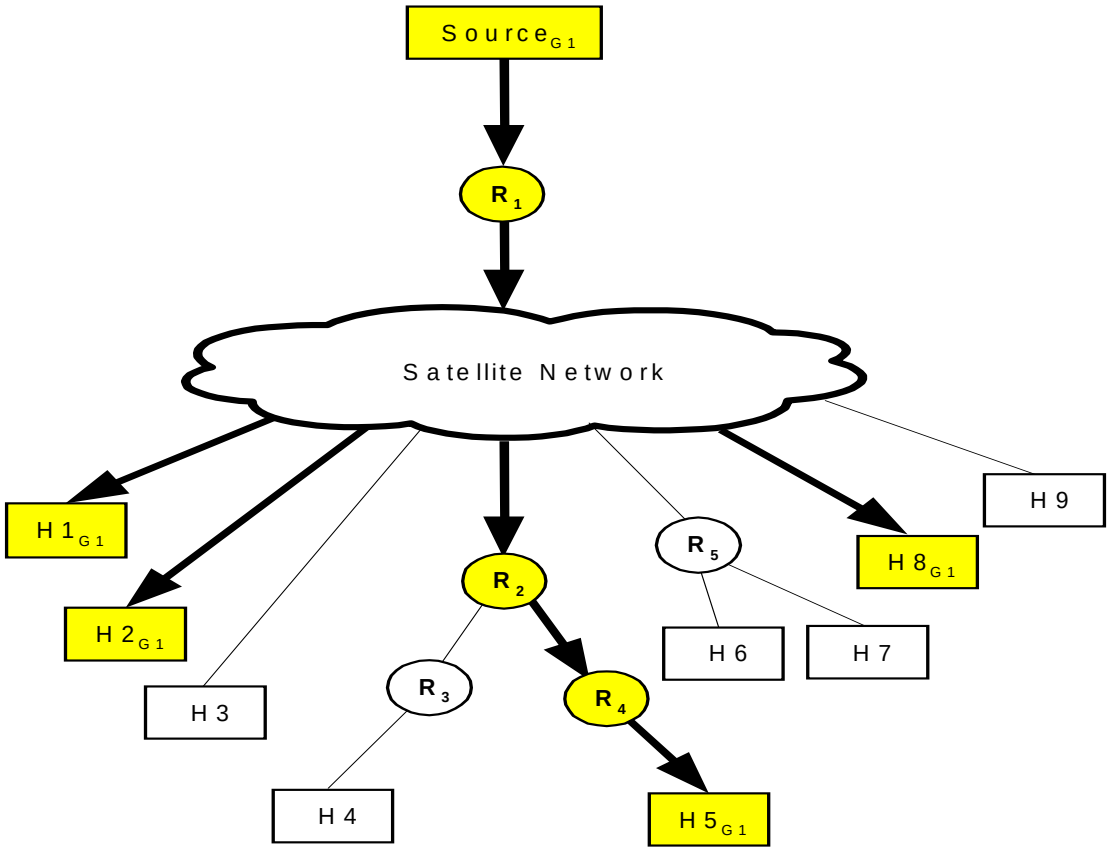

Figure 27: IP-Multicast across the satellite network (bridged)

6.5.1.3.2 Replication at IP layer

With replication at the IP layer, the satellite would be able to recognize a multicast message, and keep track of the group management. This is the most powerful solution in terms if preserving the spectrum resources, but probably also the most challenging in terms of complexity, processing and storage.

With a one-satellite solution, the satellite would be able to access all receivers. With more than one satellite it would in general be necessary to also deliver the message to other satellites.

If all satellites knew about all hosts, and where they were connected, inter-satellite traffic could be minimized, but at the expense of a solution that would require some amount of book-keeping.

The satellite could act as a multicast router. Figure 28 illustrates the multicast distribution tree for group G1. The Satellite network is acting as a layer 3 (IP) router.

The significance of this particular configuration is that the Satellite network can terminate industry standard (multicast) routing protocols and use that information to manage the space link multicast configuration (e.g. SAT MGID assignments, IP multicast group-to-MGID mapping, transport type, etc.).

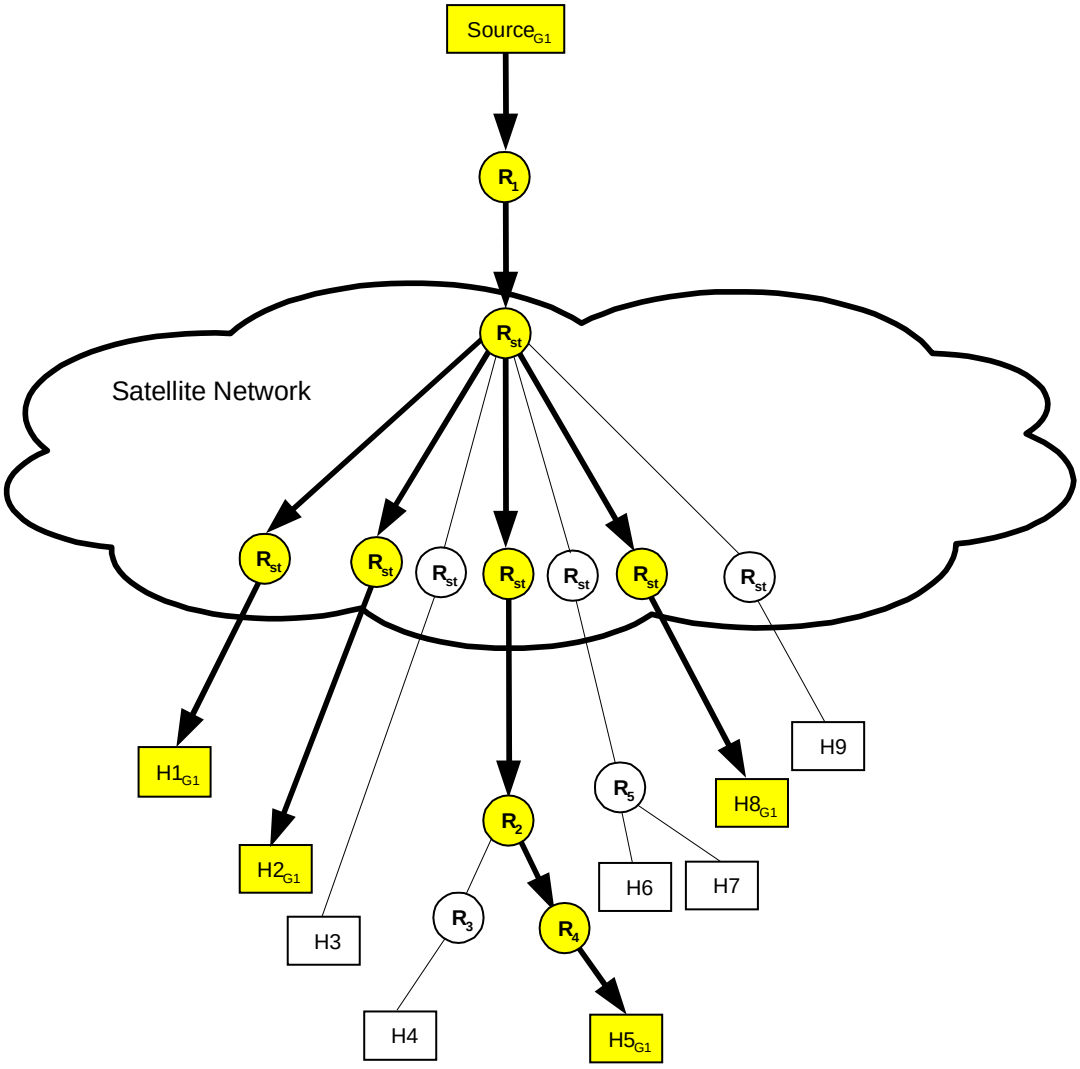

NOTE: The downstream STs (shown in the lower half of the network cloud), with hosts and/or routers locally attached, are forwarding the multicast stream onto their local LAN connections (bold arrows).

Figure 28: Replication at IP layer

6.5.2 Replication and routing summary

A key issue in BSM multicast is how to route and where to replicate the content. As has been seen from the discussion above on different satellite architectures, many factors influence routing and replication.

Replication should as a general rule be done as close to the host receiver as possible. Not looking at the architecture of the system, the following possible replication nodes exist:

- Network entry point (with a router), where content can be replicated to different gateways (and satellites, terminals and hosts).
- Gateways, where content can be replicated to different uplink beams, which again can be destined for different satellites and/or different downlink beams.
- Satellite(s), where content can be replicated to different beams.
- Terminals, where content can be replicated to different host.

These nodes form a tree, with possible branching point at the places mentioned above. Replication need not actually take place at the given nodes, either because it is not required or because it is technically impossible in some cases (like bent-pipe satellites). Still, a routing and replication table can have a general structure where no replication is handled as a (special case) 1:1 branching.

A routing and replication database should initially identify the identification of the gateways that must have a copy of the content, based upon where the destination terminals are. This table must:

1) Identify the source address.
2) Identify which terminals shall receive this source.
3) Identify where these receivers are (at the moment).
4) Identify how they are addressed (gateway, satellite, beam, carrier, timeslot).
    a) If necessary, move (group) some terminals to new carriers.
5) Identify which gateways that are to be used to reach the destination receivers.
6) Replicate the messages and send to all relevant gateways.
7) In the gateways, decide the satellites that are relevant.
    a) If satellite do not replicate decide the beams and carriers required (possibly group receivers to some carriers) and replicate the content to each beam uplink.
    b) Decide the coding and modulation.
    c) For reliable multicast, resent data when required.
8) If satellites replicate, they must replicate the data to the appropriate downlink beams, carriers and timeslots. (For bent-pipe satellites skip this point).
9) In the terminals, replicate the message to the subscribing hosts.
    a) Request any dropped packets (at level 2) for non-real-time data when reliable multicast is used.
    b) Optionally resend packets to hosts that are received ok in the terminal but that may be dropped in the host computers.

When a multicast source originates internally in the BSM system (via a satellite access terminal), then one can in general not assume that the satellite terminal will have a database identifying the other hosts in the multicast group, except if the other hosts are connected to the same receivers router. Therefore no local replication (except to locally connected hosts) will likely take place. The source must be replicated in other switches/routers, either in gateways or in satellites. Thus, the source must be tunneled to the nearest gateway (or satellite if OBP), replicated there and possibly from there on the other relevant gateways (via the NEP If required). The procedure is indicated in the routine below:

1) The terminal will identify local hosts that want local sources, and replicate the content to these IP hosts.

   a) The terminal will then forward the content either to a satellite or via the satellite to its gateway.

2) If the satellite can replicate, it will identify the source and group, and replicate to the relevant satellite terminals (beams, frequencies etc.).

   a) Next it will replicate to other satellites (if required).

   b) Then to the terrestrial segment, both for the purpose of reaching other satellites (if no ISL were present) or for reaching terrestrially connected hosts outside of the BSM system. (There may be special network exit points).

3) If forwarded via a bent-pipe satellite to a gateway the gateway will:

   a) replicate and forward to the receivers it sees; and

   b) forward the content either to a central BSM router or directly to other gateways);

   c) external hosts are addressed via external routers as mentioned above for the OBP satellite.

In the above procedures addressing schemes are needed, and routing tables could contain the addresses of the gateways, satellites, terminals and hosts.

6.5.3 Reliable multicast in BSM

It may seem a lot to request reliable multicast when using UDP as a transport layer protocol, since UDP is a connectionless protocol, unlike its counterpart Transmission Control Protocol (TCP) which is connection oriented. TCP offers a reliability we may desire in multicast, but the drawback is the overhead involved. If the packet in question makes it to the destination, as will hopefully be the case most of the time, there is no need to waste capacity for return traffic. However, if the transmission fails, we would like the receiver to be able to inform the network about this and possibly request another packet copy; similar to TCP.

It should be noted that TCP supports only data reliability; it is not suited for transport of multimedia streams, which require consistent time delivery at the receiver and only need to be semi-reliable. Thus, multimedia streaming applications need a specialized transport layer such as the Real-time Transport Protocol (RTP) for unicast as well as multicast transmissions. There are many multicast applications that require reliability rather than timeliness, similar to those unicast applications that operate over TCP, except that delivery is to several recipients rather than just one.

A potential problem with a repeat request approach is a mid-stream network outage. If a large number of receivers do not receive a packet, there will in turn be a large number of receivers requesting copies of the same packet. A packet request implosion can occur, which is a reason why TCP is not used.

A multicast protocol may (optionally) provide a delivery confirmation to ensure reliable delivery, i.e. a mechanism for receivers to inform the sender when data has been delivered. Confirmation can be either:

(i) at the application data unit level; or

(ii) at the packet level.

The main constraint on solutions is imposed if there is a need to scale to large receiver sets. For small receiver sets the design space is much less restricted. However, there are many applications for RM that do not need to scale to large numbers of receivers. For such applications, a range of solutions may be available that are not available for applications where scaling to large receiver sets is a requirement.

6.5.3.1 Totally reliable multicast (file transfer)

This can be accomplished by coding, ARQ or other means, but the important thing is that the users can rely upon the satellite system to deliver the content. If it has been detected that it has not been receiver, the content will always be retransmitted.

For BSM systems there may be a need to define a maximum loss rate (at a very low level), perhaps also limiting the number of times a packet can be re-requested. We do not want a situation where data can be re-requested for ever.

6.5.3.2 Semi-reliable multicast (audio, video)

Semi-reliable multicast will offer a very low packet loss (according to QoS contract) rate but no re-transmissions (ARQ) at the link layer. For real-time applications a retransmission will in any case not have any value.

For BSM systems there may be a need to define a maximum loss rate (at a low level), and power, coding and modulation settings could be adjusted accordingly.

6.5.3.3 Acknowledgements

Reliable multicast messages may request acknowledgements. For a satellite network the resource handling associated with this in a large group is undesirable, and it is then best avoided. For very small groups, however, it may be implemented.

Acknowledgements can be returned from a gateway if required, as long as the gateway guarantees the delivery to the destination host. This may be similar to TCP spoofing concepts.

6.5.3.4 FEC for increased reliability

Perhaps of particular interest if there is no return channel available for multicast purposes, additional FEC can be applied to the message to increase the reliability to the desired level. Messages can of course also be duplicated, but that may serve little purpose if the reason for not receiving the content has not gone away (e.g. a deep and long rain fade).

The use of FEC codes is discussed in "The Use of Forward Error Correction in Reliable Multicast", IETF draft-ietf-rmt-info-fec-03 (2003) [28]. The basic rational is the following:

- There are many ways to provide reliability for transmission protocols. A common method is to use ARQ, automatic request for retransmission. With ARQ, receivers use a back channel to the sender to send requests for retransmission of lost packets. ARQ works well for one-to-one reliable protocols, as evidenced by the pervasive success of TCP/IP.

- ARQ has also been an effective reliability tool for one-to-many reliability protocols, and in particular for some reliable IP multicast protocols. However, for one-to-very-many reliability protocols, ARQ has limitations, including the feedback implosion problem because many receivers are transmitting back to the sender, and the need for a back channel to send these requests from the receiver. Another limitation is that receivers may experience different loss patterns of packets, and thus receivers may be delayed by retransmission of packets that other receivers have lost that but they have already received. This may also cause wasteful use of bandwidth used to retransmit packets that have already been received by many of the receivers.

- In environments where ARQ is either costly or impossible because there is either a very limited capacity back channel or no back channel at all, such as satellite transmission, a Data Carousel approach (where data is repeated periodically) to reliability is sometimes used. With a Data Carousel, the sender partitions the object into equal length pieces of data, which we hereafter call source symbols, places them into packets, and then continually cycles through and sends these packets. Receivers continually receive packets until they have received a copy of each packet. Data Carousel has the advantage that it requires no back channel because there is no data that flows from receivers to the sender. However, Data Carousel also has limitations. For example, if a receiver loses a packet in one round of transmission it must wait an entire round before it has a chance to receive that packet again. This may also cause wasteful use of bandwidth, as the sender continually cycles through and transmits the packets until no receiver is missing a packet.

Forward Error Correction (FEC) codes provide a reliability method that can be used to augment or replace other reliability methods, especially for one-to-many reliability protocols such as reliable IP multicast.

6.5.4 Receiver synchronization

Some applications may require that the receiver are synchronized in such a way that all messages are received at the same time. This could for instance be the case when stock quotes are multicast or in gaming or betting situations. More generally, it could be the case for any application that would require a quick response time from the end users.

Given the focus of the BSM systems that have been treated as part of the BSM work at ETSI, such applications do not play any dominant role, however.

Further, the general situations are the following:

- In some cases only a small ratio of all multicast hosts are on any given satellite network. Presently this issue is not seen as a general requirement in the terrestrial networks.
- In other cases the satellite network itself is an autonomous network. In this case the delivery time to all hosts (and terminals) over a GEO satellite (GEO is the focus here) is fairly consistent over the set of receivers (assuming a fixed number of satellite hops). Therefore the implementations of such synchronization would in any case only play a marginal role.
- For multiple hop networks over GEO satellite it is quite possible to calculate a rough estimate for the maximum delay, which is given by the maximum number of satellite hops), and if required, some messages can be delays to compensate for this. However, a basic requirement for this has not yet been identified.

Satellite systems will however be able to report their delivery delay, and this capability should be included in the interface specification.

6.6 Multicast control functions

6.6.1 Satellite multicast routing

Satellites are inherently good for multicasting. For a spot beam satellite network the problem of how to deliver multicast messages is a mix of using broadcast channels, as in one spot beam, and routing, as in selecting spot beams and carriers. Clearly, if there are subscribers in a spot beam a multicast message needs to be sent down in that beam from the satellite. If there is more than one carrier then using only one would be best from a resource optimization point of view, but it may be that all terminals that are to receive the message do not "listen" to the same downlink carrier. Thus, the problem becomes one of resource management also.

A BSM system should not in general flood all downlink beams and all carriers with messages, unless there are relevant multicast hosts associated with these resources. However, there s a balance here, and if a clear majority of all spot beams must send the data in question, then saving one or two (or a few) spot beams will in general not save a significant amount of resources.

When considering an access network scenario, then satellite networks require a sparse mode protocol, which will spend less satellite capacity for managing the multicasting issues when the ratio of multicast subscribers to total satellite system subscribers is low. This is normally the case for satellite systems that covers both a large area and a large number of users. However, if satellites are used for edge-casting, then the content is basically to be received by all hosts (which in this case would be gateways).

If the satellite system is at the very end of the network, then it can take advantage of an optimized protocol, and the gateway(s) could be used as a multicast manager.

Multicast routing within BSM network may be static or dynamic. It should be dynamic to allow efficient use of bandwidth.

The BSM NETWORK, consisting of MEPs, "multicast exit points", hubs, satellites and terminals, should have capability to support the following IP multicast routing protocols at the BSM network edge:

- PIM
- DVMRP
- IGMP v2

- IGMP v3

PIM (sparse or dense) at the terminal allows dynamic use of the BSM network as an alternative routing way to a PIM multi-access router network connected to a satellite terminal. It should also be supported at the MEP for interconnection to other router networks.

DVMRP enables e.g. connection to the MBone. It may be available through an intermediate router converting between PIM and DVMRP.

IGMP allows dynamic host access and is required at the satellite terminal and "multicast exit point". Version 3 is the most versatile. An IGMP version 2 router has better performance for the large number of existing version 2 hosts, though a version 3 router is compatible with a version 2 host.

Dense-mode routing protocols such as DVMRP and PIM-DM, are not well suited for use over subnetworks with a large round trip delay, such as satellite networks. PIM-DM protocols rely on flooding the multicast network with packets until they receive an explicit "prune" message. This is not ideal in a BSM network because:

- The subnetwork will offer a broadband service, and there may also be a need to connect to the Internet Multicast Backbone. This backbone carries a large number of groups. Each group will (from time to time) be flooded - consuming local resource (capacity, state) - and also network resource from the upstream networks which source the data. This is the classic scaling problem that is a challenge for any multicast network but will be emphasized by some limitations of the BSM. Such limitations include limited capacity, bandwidth on demand and terminals that may be inaccessible at some point in time, large delay can result in messages timing out, etc.

- Satellite networks will from time to time have losses, and if the subnetwork has asymmetric loss or capacity then the loss of a routing "prune" message will prevent the group being pruned. This leads to a waste of spectrum capacity by continued forwarding of a group which is not required. Most satellite networks will have asymmetric capacity. If that happens then multicast messages are received by a satellite node that has zero addresses to forward it to.

- If the subnetwork round trip transit delay is appreciable, then this delays the reception of prunes. This also increases the load (i.e. may consume much more forward capacity than for a low latency subnetwork).

The use of Sparse-Mode (e.g. PIM-SM) is therefore desirable for satellite links, which have limited capacity and, can exhibit loss and have a high round trip transit delay. The use of PIM-SM with UniDirectional Link Routing (UDLR) has also been suggested as desirable. The IETF UDLR working group issued the IETF RFC 3077 [29] in March 2001. This RFC describes a link-layer tunneling mechanism (LLTM) for supporting UniDirectional Links (UDL) in the Internet.

There is on-going work to develop appropriate configuration policy for PIM-SM in hybrid (terrestrial/Satellite) networks concerning the choice for Rendezvous Points (RP) and switching from the RP-routed tree to shortest path tree. There are also issues in scaling to a large number of routing peers.

A potential problem with RPs is that they can be jammed. There is no mechanism to stop a rogue source from sending data to the RP to "jam" a broadcast. Every packet sent to a Class D address will be encapsulated and sent to the RP. Every join or leave also goes to the RP, and a RP becomes a single point of failure.

The present status is that Intra-domain routing protocols are stable, and deployed in private and local Internet. However, inter-domain routing protocols are premature, and there are concerns around traffic and policy negotiation between ISPs.

A special satellite observation is that satellite terminals that for instance include PIM routers would not be able to see other neighbour routers, as their uplinks are not visible for other terminals. Relaying these routers communications with the BSM network to all routers in the network could waste a lot of spectrum capacity, and this would not be desired. Therefore every PIM router will also be a designated router (DR). In the gateway network there could be neighbour routers however, if the network topology allowed it. Normally, one would however not expect it. A specific gateway could function as a RP for PIM-SM routers that were connected to it, and the MEP could function as a RP for all gateways. This would imply that every gateway was associated with a subnetwork. Another alternative would be to let the NCC (for instance) also be the PIM-SM RP.

With no neighbour router, the PIM "Hello" messages need in principle not be sent (but that would be non-compliant PIM-SM behaviour). The hello period could be set to maximum time.

The internal BSM network protocols must be capable of conveying multicast group membership information between the BSM network edges, as required by the supported edge protocols. The different protocol elements used within BSM network should be coordinated to optimize the use of bandwidth by reducing redundant flow of information. Use of standardized protocols within the BSM network is required in order to enable interconnection of components from different vendors. These protocols may differ from the edge protocols.

6.6.1.1 User and terminal mobility

Satellite systems can support mobile and/or nomadic users. Further, users may be mobile across fixed terminals via e.g. a SIM-card. As a consequence, a satellite system can not rely on a fixed beam location nor a fixed physical terminal if multicast is to be supported on a *user* basis. If multicast is to be supported on a terminal basis, some systems may limit the support to fixed terminals, while some may wish to support portable or mobile terminals.

The network will know where the terminal is if it is logged on, and e.g. a mapping between a fixed address and a temporary address may be implemented.

For mobility, the following multicast use-cases could be considered relevant:

1) *Terminals* that are not permanently mapped to a single beam. This covers mobile and nomadic terminals.

2) *Users* who are not permanently associated with a specific terminal.

With respect to multicast, these issues are arguments for supporting dynamic distribution trees.

6.6.2 IP multicast addressing

IP multicast provides a method of using a unique class D address to transport data to multiple destination stations using a single IP packet. As class D address range includes IP addresses between 224.0.0.0 and 239.255.255.255, concurrent multicast groups can exist over the same IP network yet remain logically independent. Applications that utilize IP multicast must establish class D address for their particular multicast session or application. Address allocation will commonly be dynamic, but static addresses may also exist for some groups.

Multicast IP addressing does not impose any new requirements in satellite systems compared with terrestrial networks, and the BSM system must be compatible with the addressing adopted for global Internet multicasting. However, address resolution, or the mapping of IP addresses into satellite layer 2 addresses, is specific to the lower layers.

Unlike IP Unicast, IP multicast addresses are not allocated to specific hosts, but instead to services, and a receiving host must identify and "listen" to one or more chosen addresses.

Multicast address allocation is an essential part of using IP multicast. Multicast addresses are an even more limited resource than unicast addresses, and must usually be allocated dynamically if they are to satisfy expected demand. (Some of the class D addresses are well known, as 224.0.1.1, which is used by NTP Network Time Protocol specified in IETF RFC 1119 [30]).

Though there are in principle 250 million multicast addresses (228) available in IPv4, these addresses are assigned globally and can get exhausted quickly as multicast usage grows. Addresses should also preferably be allocated in contiguous blocks so as to allow for multicast address aggregation in the multicast routing tables.

Multicast addresses from the allocated pool must be obtained by the originating server or service provider.

A refinement of multicast addressing has been obtained by use of a subset of the multicast address space labelled administratively scoped addresses (IETF RFC 2365 [12]) in the domain 239.0.0.0 to 239.255.255.255. This prevents the forwarding of IP multicast packets outside administratively restricted domains. This mechanism is much more efficient than the current use of TTL-scoped addressing (using small TTL values restricts the distribution of multicast packets when large TTL decrements are applied in border routers), using the TTL field in the IP header.

Typical MBone (Multicast Backbone) usage has been to engineer TTL thresholds that confine traffic to some administratively defined topological region. The basic forwarding rule for interfaces with configured TTL thresholds is that a packet is not forwarded across the interface unless its remaining TTL is greater than the threshold.

Administratively scoped addresses will enable multicast technology to be used for communication among small user groups (e.g. videoconferencing) without spreading the associated state information all over the Internet (which would be hard to justify, regarding the fairly limited savings in bandwidth).

6.6.3 MAC layer multicast addressing

The address space in the L2 (MAC) layer may vary between different satellite systems. With a large number of spot-beams and many multicast groups in addition to other types of traffic the address space of the satellite system must be considered.

In DVB and DVB-RCS it is for example common to use (Packet ID) PID for addressing different terminals. The DVB MPEG TS PID is a 13-bit field (0x1fff, or 8 191 channel) indicating the type of data stored in the packet payload. A PID defines a unidirectional broadcast channel. There is no specified standard for mapping of IP addresses to PIDs. In a multiple spot-beam system with switching capability at least one PID must be assigned to every beam. Filtering of information could be performed at the IP layer in the satellite access terminals, at the expense of using data processing. If the carriers were much broader than today and higher order modulation was used, then both these factors would require additional processing power in the terminals.

For multicasting, assigning multicast addresses to special PIDs is possible, but there is no chance to map all IP addresses to different PIDs. One can send all IP content in one PID or assign particular groups of IP addresses to separate PIDs and do filtering in the receivers. As some receivers may not able to support a large number of parallel PIDs, such mappings should be used with care, not requiring too many PIDs.

For DVB delivery of multicast, (IETF Internet draft "Requirements for IP transport over DVB") issues are mentioned to include:

- Mapping IP multicast groups to the underlying MPEG-2 TS logical channel (PID) and the MPEG-2 TS Multiplex.
- Provide signalling information to allow a receiver to locate an IP multicast flow within an MPEG-2 TS Multiplex.

The draft also mentions that appropriate procedures need to be specified to identify the correct action when the same multicast group is available on separate TS logical channels. This could arise when different end hosts act as senders to contribute IP datagrams with the same IP group destination address. Another different case arises when a receiver may potentially receive more than one copy of the same packet. In some cases, these may be sent in different TS logical channels, or even different TS Multiplexes. The primary goal of multicast support will be efficient filtering of IP-multicast packets by the receiver, and the mapping of IPv4 and IPv6 multicast addresses onto the associated PID value and TS Multiplex. The design should permit a large number of active multicast groups, and should minimize the processing load at the receiver when filtering and forwarding IP multicast packets.

6.6.4 Satellite multicast group management

There are several options for group management depending on the requirements in question. For a satellite network where the spectrum resources are limited and should be used carefully, it is of importance to take this into account when designing satellite group management system.

A multicast router *periodically* sends queries to all hosts participating in IP multicast on the special 224.0.0.1 all-hosts group. Each relevant host sets a random timer for each group it is a member of. When the timer expires, it sends a report message on that group multicast address. Each host that gets a report message for a group cancels its local timer for that group. When a host joins a group it announces that on the group multicast address. The Max Response Time field in IGMP is used only in Membership Query messages. It specifies the maximum allowed time before sending a responding report in units of 1/10 second. In all other messages, it is set to zero by the sender and ignored by receivers. Varying this setting allows IGMPv2/3 routers to tune the "leave latency" (the time between the moment the last host leaves a group and when the routing protocol is notified that there are no more members). It also allows tuning of the burstiness of IGMP traffic on a subnet.

A satellite network configuration has a number of characteristics that can be exploited in order to simplify the multicast group management. The architecture is fairly stable. New satellites are not added every day, neither are new gateways. Some degree of dynamic capabilities may still be desired to allow for rerouting in case an element fails (like changing to another gateway in case of deep rain fades). New satellite terminals will generally be added relatively frequently, but they will be added in a structured manner and usually the need to be authorized to operate before they will be allowed to handle traffic. It may be acceptable to handle network changes a few times a day, at least with respect to multicast capabilities. All changes will under normal circumstances be changes to the terminal segment, and all such changes will have the same character; either the installation (or removal) of a satellite terminal. A terminal will have one or a few computers connected to it via local router functionality (if more than one). Business solutions will include in contrast a "full scale" stand-alone router.

Group management for joining and leaving groups for users could be done using a standard API to the user, and not needing some special interface or protocol. Users would need a service discovery mechanism, which could be a common broadcast channel or even a web page, supporting multiple platforms and a sensible protocol that scales and thinks about IP networks. How should groups normally be joined and left? That could be a service provider issue, but a first stage solution could be via a web-page hosted by the service provider. By expressing interest in particular groups the host coordinates could be entered directly in a database that in turn could interact with the routing table. If local routing was needed at the terminal side, then the terminal would have to be told which computer host that subscribed. Another option is that the groups are "proxy cached" by the terminal, and that joining and leaving is via a local web-based (man machine) interface. The terminal would then inform the gateway/BSM network about the source-group activity change.

There is also the possibility to include a simple IGMP router functionality in the terminal, that "looks and feels" as an IGMP router to the host computer (which for instance is running windows XP). In this case the terminal would (could) spoof the functionality with respect to timers and queries to conserve satellite bandwidth. When there was a change in (S,G) settings this would be communicated to the network. Thus the terminal would have an IP/IGMP interworking function, and the group management over the satellite segment could be BSM specific.

In a satellite network the multicast group management messages sent using various timers are less needed, due to the more static nature of the network. Further, a satellite network does not want to spend spectrum capacity on a service that carriers little or no new information (entropy). The philosophy being that it is better to signal the changes when there are some, than to keep asking if there are any. The gateways connecting to the outside world would however be required to respond normally to PIM timers, and the satellite terminals would need to request IGMP information form the hosts. But with some basic filtering, the terminals should not need to transmit "no change" over the satellite.

The gateway or the satellite itself would interact with the terminals and the (S,G) database in the network would be updated. The network would also decide the routing from the desired source to the terminal (the subscribing host itself need not be visible, as the local routing is handled by the terminal). The routing could be static or dynamic. A given terminal may be served by more than one gateway, and if the constellation is not GEO then different satellites will be involved. However, we will focus on GEO satellites here, in which case a fixed terminal will always be in the same spotbeam, but not necessarily on the same carrier, and not necessarily served always by the same gateway (but by the same satellite).

Portable terminals would require updating the routing database when logging on if they have moved. Mobile terminals (not in focus here) would in any case be tracked and this information would be used for updating multicast routing databases.

The BSM network would appear as a subnetwork to an external terrestrial router, and multicast groups could be managed via a PIM-SM (or PIM-SSM) router interface. It would be essential that the BSM network would not let itself be jammed for instance by letting unauthorized and possible hostile (hacker) sources start spending satellite capacity. Therefore, a proxy/spoofing mechanism in the BSM network would be beneficial also in this respect. Large, corporate or campus networks would require local PIM routers and "PIM-SM" enabled corporate terminals.

The different options for group management discussed here are illustrated below. To summarize, the consumer terminals would use IGMP to the hosts, but spoof the protocol to only send the changes. The corporate networking solution may need to interact with PIM-SM, and the PIM-SM timers should be set to minimize over-the-air traffic. But if there are few PIM-SM setups then the problem of wasting capacity is small. The figure shows OBP satellites with IGMP capability, which can be one option.

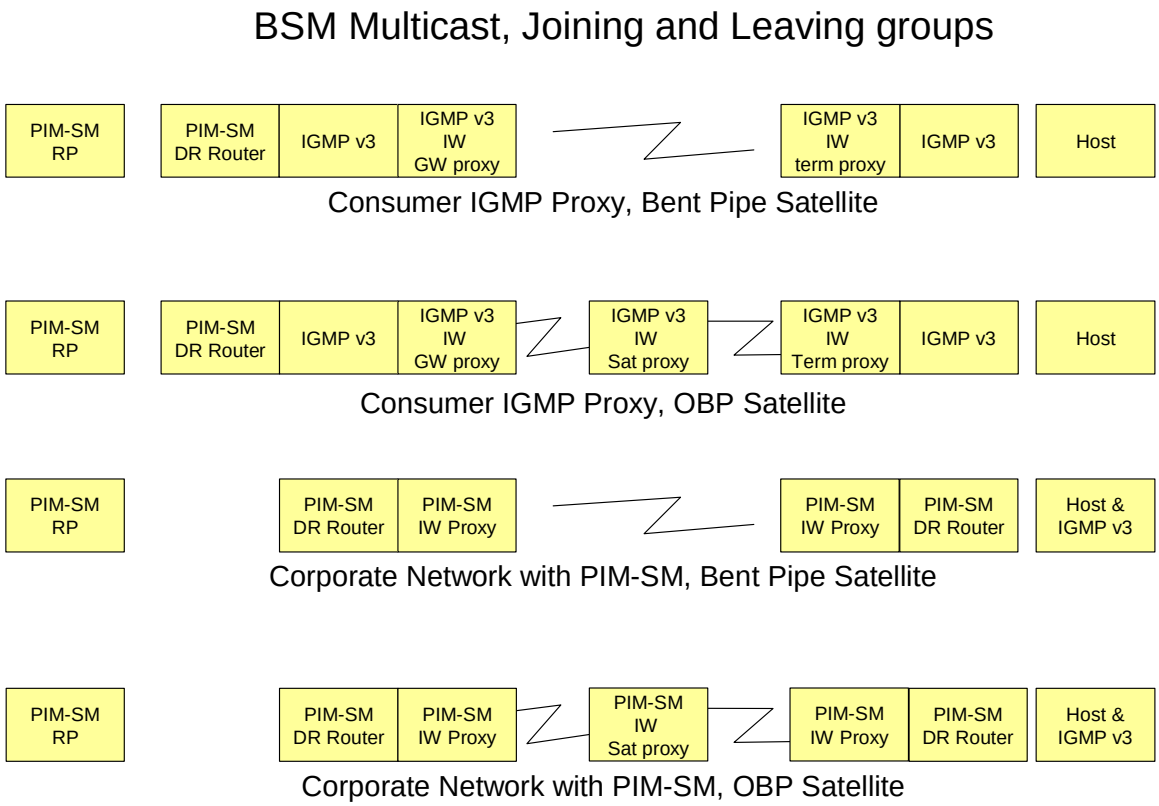

Figure 29: Alternative join/leave configurations

6.7 Multicast management functions

6.7.1 Capacity requirements and resource management

A RRM scheme will have to allocate terminals that shall receive the same multicast content to the same carrier in a spot beam, if several are present. The RRM system will also have to make sure the terminals either listen to the timeslots where the multicast content is sent, unless they listen to all content.

6.7.2 Traffic management

Multicast group management and link layer protocols requiring retransmissions and using acknowledgements may be extremely challenging to handle if the concept is not properly designed.

Traffic management and congestion control on the forward channel from a gateway may not be of the highest importance, as the gateway itself normally handles a large portion of the traffic. We can also assume that the terrestrial connection in to the gateway is of a high data-rate. We can also assume that the data input rate to the gateways (or to the multicast replication nodes) will in general be less than the total capacity of the gateway to relay this to the satellite. If however, this is not the case, then either the data must be rerouted to another gateway with available uplink capacity or packets must be dropped. The gateway will interface to the core network via common IP protocols and commonly applied terrestrial congestion control (such as slow start) can be applied.

The scenario is however quite different on the return channel from the receiving hosts. The importance of, and the need for congestion control mechanisms are essential. This is in particular true if the hosts are required to frequently acknowledge data or respond to administrative messages. Return traffic may in such cases come from a large group of receivers, and for satellite systems it may be specifically large due to the large coverage area a satellite has. In particular congestion control may be a challenge for bursty sources and multiparty conference situations (that in turn would use multicast to relay conference data to all parties). Return acknowledgements and other responses could, if not designed to do otherwise, come in bursts just after a message was relayed, and overload the system. If congested the gateways would not receive acknowledgements and could in turn start retransmitting data in some case making the situation even worse.

A key element in a BSM multicast design will therefore be to demand as little return traffic as possible yet be tolerant to some amount of congestion. Sparse mode protocols are therefore an obvious choice. ARQ should be used with care only when required (e.g. not for real-time audio and video applications).

6.7.3 Physical and Link Layer issues

It is on the physical layer the satellite multicast may differ from other technologies for delivering multicast content in that it truly is a shared physical resource. Granted, it is so for a number of other physical layers, such a radio networks and shared cable networks. However, none of these can boast the same coverage area as a satellite cell, and thus in an area as large as a satellite cell other technologies would require a number of different or similar networks. Specifically, one satellite can cover about 1/3 of the face of the earth. The coverage area is a key factor on the positive side when considering multicast capabilities. However, the large coverage of a GEO satellite comes with the delay of around 250 ms (return). This will influence some protocols, timers etc.

Multimedia satellite systems will primarily operate in the Ku and/or the Ka band in the near to medium future. Some systems will offer broadband services in lower frequency bands, but in the higher bands commercial services are not expected for the next 5-10 years (see TR 101 374-2 [2]). The Ka band, in particular, can be subject to rain fading, leading to either a reduced channel capacity or to an outage. Fades and outages will be of regional character, and even in the same beam some satellite terminals may fade while others do not with conditions otherwise the same except for the rain. This may have an impact on multicasting, and given a large and diverse enough region there will always be a finite probability that some terminal will be unreachable at any given time. The question then becomes one of how to handle the multicast in these cases. Buffering and retransmitting is one alternative. Dropping the data is another. The specific reaction needs to be decided based upon potential QoS settings and the type of data. Real time voice and video makes little or no sense to buffer, while offline file downloads do.

Most BSM systems will use a shared or common TDM carrier from the satellite to the terminals (downlink). In a spot beam there may be more than one TDM, but on the other hand there may be less than one in some sparsely populated regions. In the latter case there may be a scanning spot beam that only lights up a region for a short while if there is traffic to send. Terminals may not listen to everything that is sent on a TDM, but may only process parts of the carrier where it expects traffic. For battery-operated devices (low power consumption) this is in particular true. Thus a resource management system will have to work to make sure all terminals that shall receive the same multicast traffic listen to the same timeslots.

A return channel is typically a MF/TDMA channel, i.e. there are a number of more narrow carriers that are shared by a number of terminals. Thus while it would be normal for all terminals in a region to listen to one and the same downlink carrier, it would also be expected that a number of different return channels would be used by the same terminals. Depending on how the terminals are allocated resources on the return link, there may be a chance that the return link congests unless traffic is shaped properly. Thus the Multiple Access Control must be suited for the type of traffic considered. For multicast return traffic this may be different than for a unicast return link.

When multicast content enters the satellite network, some of the destination terminals may already have been logged off from the network, or be in a form of standby mode. They may or may not be able to receive content, but when not online they may have no return link resources allocated. Terminals may have to go through a normal log-on procedure (taking several seconds) before they are able to send on their return link.

6.7.4 Performance and QoS issues

In the long term QoS must be supported by BSM multicast systems. In the short term QoS for multicast is not used (as of March 2003).

Multicast QoS classification can be seen as:

- Best effort, with no user specifications and no network guarantees. A best-effort routing protocol (e.g. PIM-SM) sets up a layer-3 routing tree.
- Flow based (like IntServ), where individual receivers explicitly specify QoS requirements and the network delivers per-user guarantee. Guaranteed service can be using RSVP, where resource-reservation is performed at a layer above the routing. RSVP signalling creates a reservation tree by discovering and reserving resources along the routing tree.
- Aggregated flows (like DiffServ), where every receiver has implicit and class-based quality expectations, and the network attempts to guarantee multicast QoS to every class in an aggregated basis.

Some relevant receiver specified parameters could be:

- Specification of data-rate.
- Requirements for an end-to-end delay-bound.
- Requirements on end-to-end loss-bound.

Settings can be heterogeneous, where QoS is independently defined between every sender-receiver pair within a multicast group, or settings can be dynamic, where the QoS specification for a receiver can change dynamically.

The IntServ architecture that IETF started work on in 1993, is based on per-flow resource reservation, where the goal of proving real-time services simultaneously with traditional non-real-time services in a shared IP network. The IntServ model requires that source and destination hosts exchange RSVP signalling messages to establish packet classification and forwarding state at each node along the path between them. The resource reservation setup protocol (RSVP) was part of IntServ, and allows individual applications to request resources from routers and then installs per-flow state along the path of the packet flow. Guaranteed service models were introduced, providing firm assurances (through strict admission control, bandwidth allocation, and fair queuing) for applications that require guaranteed bandwidth and delay. Flow specifications provide a syntax that allows applications to specify their specific resource requirements. In September 2001, IETF RFC 3175 defined procedures that allow a single RSVP reservation to aggregate other RSVP reservations across a large IP network. Recent additions to the IntServ architecture enhance the scalability of RSVP by allowing it to aggregate resource reservations and to potentially play a more significant role in large IP networks.

The DiffServ architecture is an alternative to IntServ that provides the scalability for deployment in large IP networks when attempting to offer better than best-effort service. The complete DiffServ architecture, where work stared in 1998, is defined in IETF RFC 2475, is based on a relatively simple model, whereby traffic that enters a network is first classified and then possibly conditioned at the edges of the network. The goal was to create relatively simple and coarse methods of providing differentiated classes of service for Internet traffic. A combination of traffic conditioning (policing and shaping) at the edges of the network, packet marking at the edges of the network, local per-class forwarding behaviours in the interior of the network, and adequate network provisioning allow the DiffServ model to support scalable service discrimination across a common IP infrastructure. Routers require buffer memory to absorb temporary bursts, so that packets are not immediately dropped when congestion occurs. Sustained congestion causes packets to be dropped. The fundamental idea behind the DiffServ model is that the deployment of multiple queues on a port allows a router to service certain traffic before other types of traffic, and thus isolate congestion to a subset of a router's queues. The deployment of multiple queues on a port allows some queues to experience congestion while other queues do not. This approach is based on the assumption that it is acceptable for less important traffic not to have access to network resources during periods of congestion. The primary problem with this assumption is that sustained congestion for this class will result in a poor experience for users.

MultiProtocol Label Switching (MPLS) supports the convergence of two fundamentally different approaches to data networking (datagram and virtual circuit) in a seamless and fully integrated manner. MPLS traffic engineering reduces congestion and optimizes the use of existing network resources by allowing you to carefully manage the distribution of traffic across your network.

The use of integrated services over DiffServ networks is significantly more complex for multicast sessions than for unicast sessions. With respect to a multicast connection, each participating region has a single ingress router and zero, one or several egress routers. The difficulties of multicast are associated with DiffServ regions that contain several egress routers. There are some major problems associated with heterogeneous multicast trees containing branch points within the DiffServ region, i.e. multicast trees where the level of resource requirement is not uniform among all receivers, i.e. two receivers requesting the same source with different QoS.

For further information and discussion on QoS, please refer to TR 102 157.

6.8 Security in satellite multicast

There is work in progress on security in the ETSI BSM WG, and basic BSM security will be discussed there.

Multicast applications are no different than unicast applications with respect to their need for security, and they require the same basic security services: user authentication, data integrity, data privacy and user privacy (anonymity). However, enabling security for multicast applications is even more of a challenge than for unicast. Having multiple receivers makes a difference, as does their heterogeneity and the dynamic nature of multicast group memberships.

Multicast security requirements (from IETF RFC 3170) can include any combination of the following services:

- Limiting Senders - Controlling who can send to group addresses.
- Limiting Receivers - Controlling who can receive.
- Limiting Access - Controlling who can "read" multicast content either by encrypting content or limiting receivers (which is not possible yet).
- Verifying Content - Ensuring that data originated from an authenticated sender and was not altered en route.
- Protecting Receiver Privacy - Controlling whether sender(s) or other receivers know receiver identity.
- Firewall Traversal - Proxying outgoing "join" requests through firewalls, allowing incoming or outgoing traffic through, and (possibly) authenticating receivers for filtering purposes and security.

The challenge of security in satellite environments is considered to be one of the main issues to solve prior to widespread deployment of satellite IP multicast and multimedia applications. The main problem is that eavesdropping and active intrusion is easier than in terrestrial fixed or mobile networks because of the broadcast nature of satellites. In addition, satellite channels experience long delays and high bit error rates, which may cause the loss of security synchronization. This demands a careful evaluation of encryption systems to prevent Quality of Service (QoS) degradation because of security processing. Also the number of members in a multicast group can be very large and change dynamically.

- Key distribution. How do we update keys when people join and leave groups, in a dynamic and well-populated environment with a minimum use of satellite resources?
- What are the dynamics of the group?

7 Other multicast standards work

This clause identifies other relevant standardization work related to IP multicasting over satellite networks. The majority of the work has been done by the IETF, but the ITU is also doing relevant and interesting work.

Not noted here, but also of possible interest, are the numerous projects, both national and international, that are working on satellite multicast. Several of these projects, specifically EU projects, have indicated in their terms of reference that they shall interact with standards bodies. Some of these can be considered as standards input bodies.

7.1 DVB

It is not known that the DVB project is working specifically on multicast work, but certainly there has been work on multicast over DVB.

DVB-S with MPE is a layer 2 multicast capable protocol. One of the addressing options is to adapt to the multicast MAC address scheme specified by IETF RFC 1112 [7]. This mapping uses IETF RFC 1112 [7] to give the rules for configuring the filtering at layer 2 in the satellite terminal receiver.

No dynamic multicast signalling has been defined for DVB-RCS, but DVB-RCS could use a multicast protocol resembling PIM-SM and IGMP v3.

7.2 IETF multicasting standardization work

7.2.1 TCP/IP network model

Although the OSI model is widely used and often cited as the standard, TCP/IP protocol has been used by most Unix workstation vendors. TCP/IP is designed around a simple four-layer scheme. It does omit some features found under the OSI model. Also it combines the features of some adjacent OSI layers and splits other layers apart. The four network layers defined by TCP/IP model are as follows.

- Layer 1 - Link, This layer defines the network hardware and device drivers.
- Layer 2 - Network, This layer is used for basic communication, addressing and routing. TCP/IP uses IP and ICMP protocols at the network layer.
- Layer 3 - Transport, Handles communication among programs on a network. TCP and UDP falls within this layer.
- Layer 4 - Application, End-user applications reside at this layer. Commonly used applications include NFS, DNS, arp, rlogin, talk, ftp, ntp and traceroute.

7.2.2 Relevant IETF Working Groups

This clause lists some of the most relevant IETF workgroups. Please see the IETF website for detailed charters and work item status.

- Inter-Domain Multicast Remnants (IDMR) deals with IGMP, and also with DVMRP.
- Protocol Independent Multicast remnants (PIM) deals with PIM and is chartered to standardize and promote the Protocol Independent Multicast Version 2, Sparse Mode and Dense Mode, as a scalable, efficient and robust multicast routing protocol, capable of supporting thousands of groups, different types of multicast applications, and all major underlying layer-2 subnetwork technologies.
- Audio/Video Transport (AVT) deals with streaming audio, video and multimedia using the Real Time Protocol (RTP).
- Reliable Multicast Transport (RMT) deals with reliable transport proposals.
- Multiparty Multimedia Session Control (MMUSIC) chartered develop protocols to support Internet teleconferencing sessions.
- Session Initiation Protocol (SIP) is chartered to continue the development of SIP, currently specified as proposed standard IETF RFC 2543. SIP is a text-based protocol, similar to HTTP and SMTP, for initiating interactive communication sessions between users. Such sessions include voice, video, chat, interactive games, and virtual reality.
- MBone Deployment (MBoneD) deals with issues of MBone deployment.
- Multicast Address Allocation (MALLOC) deals with one means of actively managing the multicast address space.
- Source Specific Multicast (SSM).
- Large Scale Multicast Applications (LSMA).
- Service Localization Protocol (SVRLOC).
- Resource Reservation Setup Protocol (RSVP).
- MAGMA, Multicast and Anycast Group Membership. This working group will be responsible for developing the functionalities required for group membership reporting and other related actions.
- UDLR, UniDirectional Link Routing, for supporting UniDirectional Links (UDL) in the Internet.
- MSEC, Multicast Security. The purpose of the MSEC WG is to standardize protocols for securing group communication over internets, and in particular over the global Internet. Initial efforts will focus on scalable solutions for groups with a single source and a very large number of recipients.
- Border Gateway Multicast Protocol (BGMP), the supposed long term solution for multicast inter-domain routing.
- Source Specific Multicasting (SSM) deals with a proposal for removing complexity from PIM-SM and thereby making multicast more practical.

There is also an IRTF Working Group on secure multicast.

7.2.3 Multicast or satellite related RFC documents

A search in the RFC database for satellite and multicast gives a large number of hits. Some of those considered most relevant are listed below.

Best current practice

- draft-ietf-pilc-link-design-12.txt (July 2002)
  - This document provides advice to the designers of digital communication equipment, link-layer protocols and packet-switched subnetworks (collectively referred to as subnetworks) who wish to support the Internet protocols but who may be unfamiliar with Internet architecture and the implications of their design choices on the performance and efficiency of the Internet.
- IETF RFC 2365 [12]: Administratively Scoped IP multicast (1998)
  - This document defines the "administratively scoped IPv4 multicast space" to be the range 239.0.0.0 to 239.255.255.255. In addition, it describes a simple set of semantics for the implementation of Administratively Scoped IP multicast. Finally, it provides a mapping between the IPv6 multicast address classes (IETF RFC 1884 [13]) and IPv4 multicast address classes.
- IETF RFC 3171: IANA Guidelines for IPv4 Multicast Address Assignments
  - This memo provides guidance for the Internet Assigned Numbers Authority (IANA) in assigning IPv4 multicast addresses.

Experimental

- IETF RFC 2934: Protocol Independent Multicast MIB for IPv4 (2000)
  - This memo defines a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. In particular, it describes managed objects used for managing the Protocol Independent Multicast (PIM) protocol for IPv4.
- IETF RFC 2909: The Multicast Address-Set Claim (MASC) Protocol (2000)
  - This document describes the Multicast Address-Set Claim (MASC) protocol which can be used for inter-domain multicast address set allocation. MASC is used by a node (typically a router) to claim and allocate one or more address prefixes to that node's domain. While a domain does not necessarily need to allocate an address set for hosts in that domain to be able to allocate group addresses, allocating an address set to the domain does ensure that inter-domain group-specific distribution trees will be locally-rooted, and that traffic will be sent outside the domain only when and where external receivers exist.

Informational

- IETF RFC 3353 [18]: Overview of IP Multicast in a Multi-Protocol Label Switching (MPLS) Environment (2002)
  - This document offers a framework for IP multicast deployment in an MPLS environment. Issues arising when MPLS techniques are applied to IP multicast are overviewed. The pros and cons of existing IP multicast routing protocols in the context of MPLS are described and the relation to the different trigger methods and label distribution modes are discussed. The consequences of various layer 2 (L2) technologies are listed. Both point-to-point and multi-access networks are considered.

- IETF RFC 3048: Reliable Multicast Transport Building Blocks for One-to-Many Bulk-Data Transfer (2001)

  - This document describes a framework for the standardization of bulk-data reliable multicast transport. It builds upon the experience gained during the deployment of several classes of contemporary reliable multicast transport, and attempts to pull out the commonalities between these classes of protocols into a number of building blocks. To that end, this document recommends that certain components that are common to multiple protocol classes be standardized as "building blocks". The remaining parts of the protocols, consisting of highly protocol specific, tightly intertwined functions, shall be designated as "protocol cores". Thus, each protocol can then be constructed by merging a "protocol core" with a number of "building blocks" which can be re-used across multiple protocols.

    - Network Topologies. A protocol must not break the network when deployed in the full Internet. However, we recognize that intranets will be where the first wave of deployments happen, and so are also very important to support. Thus, support for **satellite networks** (including those with terrestrial return paths or no return paths at all) is encouraged, but not required.

- IETF RFC 2908: The Internet Multicast Address Allocation Architecture (2000)

  - This document proposes a Multicast address ALLOCation architecture (MALLOC) for the Internet. The architecture is modular with three layers, comprising a host-server mechanism, an intra-domain server-server coordination mechanism, and an inter-domain mechanism.

- IETF RFC 2887 [6]: The Reliable Multicast Design Space for Bulk Data Transfer (2000)

  - The design space for reliable multicast is rich, with many possible solutions having been devised. However, application requirements serve to constrain this design space to a relatively small solution space. This document provides an overview of the design space and the ways in which application constraints affect possible solutions.

- IETF RFC 2771: An Abstract API for Multicast Address Allocation (2000)

  - This document describes the "abstract service interface" for the dynamic multicast address allocation service, as seen by applications. While it does not describe a concrete API (i.e. for a specific programming language), it describes - in abstract terms - the semantics of this service, including the guarantees that it makes to applications.

- IETF RFC 2729: Taxonomy of Communications Requirements for Large-Scale Multicast Applications (1999)

  - The intention of this memo is to define a classification system for the communication requirements of any large-scale multicast application (LSMA). It is very unlikely one protocol can achieve a compromise between the diverse requirements of all the parties involved in any LSMA. It is therefore necessary to understand the worst-case scenarios in order to minimize the range of protocols needed. Dynamic protocol adaptation is likely to be necessary which will require logic to map particular combinations of requirements to particular mechanisms. Standardizing the way that applications define their requirements is a necessary step towards this. Classification is a first step towards standardization.

- IETF RFC 2627 [15]: Key Management for Multicast: Issues and Architectures (1999)

  - This report contains a discussion of the difficult problem of key management for multicast communication sessions. It focuses on two main areas of concern with respect to key management, which are, initializing the multicast group with a common net key and rekeying the multicast group. A reek may be necessary upon the compromise of a user or for other reasons (e.g. periodic reek). In particular, this report identifies a technique which allows for secure compromise recovery, while also being robust against collusion of excluded users. This is one important feature of multicast key management which has not been addressed in detail by most other multicast key management proposals. The benefits of this proposed technique are that it minimizes the number of transmissions required to rekey the multicast group and it imposes minimal storage requirements on the multicast group.

- IETF RFC 2588: IP Multicast and Firewalls (1999)

  - In this document, we discuss the issues surrounding the traversal of IP multicast traffic across a firewall, and describe possible ways in which a firewall can implement and control this traversal. We also explain why some firewall mechanisms - such as SOCKS - that were designed specifically for unicast traffic, are less appropriate for multicast.

- IETF RFC 2502: Limitations of Internet Protocol Suite for Distributed Simulation the Large Multicast Environment (1999)

  - The Large-Scale Multicast Applications (LSMA) working group was chartered to produce documents aimed at a consensus based development of the Internet protocols to support large scale multicast applications including real-time distributed simulation. This memo defines services that LSMA has found to be required, and aspects of the Internet protocols that LSMA has found to need further development in order to meet these requirements.

- IETF RFC 2490: A Simulation Model for IP Multicast with RSVP (1999)

  - This document describes a detailed model of IPv4 multicast with RSVP that has been developed using the OPNET simulation package, with protocol procedures defined in the C language. The model was developed to allow investigation of performance constraints on routing but should have wide applicability in the Internet multicast/resource reservation community.

- IETF RFC 2432: Technology for IP Multicast Benchmarking (1998)

  - The purpose of this document is to define terminology specific to the benchmarking of multicast IP forwarding devices. It builds upon the tenets set forth in IETF RFC 1242, IETF RFC 2285, and other IETF Benchmarking Methodology Working Group (BMWG) efforts. This document seeks to extend these efforts to the multicast paradigm.

- IETF RFC 2375: IPv6 Multicast Address Assignments (1998)

  - This document defines the initial assignment of IPv6 multicast addresses. It is based on the "IP Version 6 Addressing Architecture" and current IPv4 multicast address assignment. It adapts the IPv4 assignments that are relevant to IPv6 assignments. IPv4 assignments that were not relevant were not converted into IPv6 assignments. Comments are solicited on this conversion.

- IETF RFC 1458 [14]: Requirements for Multicast Protocols (1993)

  - Multicast protocols have been developed over the past several years to address issues of group communication. Experience has demonstrated that current protocols do not address all of the requirements of multicast applications. This memo discusses some of these unresolved issues, and provides a high-level design for a new multicast transport protocol, group address and membership authority, and modifications to existing routing protocols.

- IETF RFC 3170: IP Multicast Applications: Challenges and Solutions (2001)

  - This document describes the challenges involved with designing and implementing multicast applications. It is an introductory guide for application developers that highlights the unique considerations of multicast applications as compared to unicast applications.

Standards track

- IETF RFC 2933 [37]: Internet Group Management Protocol MIB (2000)

  - This memo defines a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. In particular, it describes objects used for managing the Internet Group Management Protocol (IGMP).

- IETF RFC 2932: IPv4 Multicast Routing MIB (2000)

  - This memo defines a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. In particular, it describes managed objects used for managing IP Multicast Routing for IPv4, independent of the specific multicast routing protocol in use.

- IETF RFC 2710: Multicast Listener Discovery (MLD) for IPv6 (1999)
  - This document specifies the protocol used by an IPv6 router to discover the presence of multicast listeners (that is, nodes wishing to receive multicast packets) on its directly attached links, and to discover specifically which multicast addresses are of interest to those neighbouring nodes. This protocol is referred to as Multicast Listener Discovery or MLD. MLD is derived from version 2 of IPv4's Internet Group Management Protocol, IGMPv2. One important difference to note is that MLD uses ICMPv6 (IP Protocol 58) message types, rather than IGMP (IP Protocol 2) message types.
- IETF RFC 2417: Definitions of Managed Objects for Multicast over UNI 3.0/3.1 Based ATM Networks (Obsoleted RFC 2366) (1998)
  - This memo defines a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. In particular, it describes managed objects for IP hosts and routers that use a Multicast Address Resolution Server (MARS) to support IP multicast over ATM, as described in "Support for Multicast over UNI 3.0/3.1 based ATM Networks".
- IETF RFC 2366: Definitions of Managed Objects for Multicast over UNI 3.0/3.1 Based ATM Networks (Obsoleted by IETF 2417) (1998)
  - This memo defines a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. In particular, it describes managed objects for IP hosts and routers that use a Multicast Address Resolution Server (MARS) to support IP multicast over ATM, as described in "Support for Multicast over UNI 3.0/3.1 based ATM Networks".
- IETF RFC 2327: Session Description Protocol (1997)
  - This document defines the Session Description Protocol, SDP. SDP is intended for describing multimedia sessions for the purposes of session announcement, session invitation, and other forms of multimedia session initiation.
- IETF RFC 3307 [16]: Allocation Guidelines for IPv6 Multicast Addresses (August 2002) (Proposed Standard)
  - This document specifies guidelines that must be implemented by any entity responsible for allocating IPv6 multicast addresses. This includes, but is not limited to, any documents or entities wishing to assign permanent IPv6 multicast addresses, allocate dynamic IPv6 multicast addresses, and define permanent IPv6 multicast group identifiers. The purpose of these guidelines is to reduce the probability of IPv6 multicast address collision, not only at the IPv6 layer, but also at the link-layer of media that encode portions of the IP layer address into the MAC layer address.
- IETF RFC 3306: Unicast-Prefix-based IPv6 Multicast Addresses (August 2002)
  - This specification defines an extension to the multicast addressing architecture of the IP Version 6 protocol. The extension presented in this document allows for unicast-prefix-based allocation of multicast addresses. By delegating multicast addresses at the same time as unicast prefixes, network operators will be able to identify their multicast addresses without needing to run an inter-domain allocation protocol.
  - The current IPv4 multicast address allocation architecture (IETF RFC 2908) is based on a multi-layered, multi-protocol system. The goal of this proposal is to reduce the number of protocols that need to be deployed in order to get dynamic multicast address allocation.

7.3 ITU multicasting standardization work

In a March 2002 presentation titled "Multicast Delivery of Broadband Multimedia Applications and Services", Seok Joo Koh, Editor of X.606 in Q.8/17, informs about ITU work related to multicasting. Claims for status:

- Intra-domain routing protocols are stable, and deployed in private and local Internet.
- Inter-domain routing protocols are premature, and there are traffic/policy negotiation issues between ISPs.

The ITU-T Q.8/17 is working on ECTP (Enhanced Communications Transport Protocol), and there is also focus on QoS Management, QoS negotiations/monitoring/maintenance issues. There is also work on a group management protocol including:

- Session Management.
- Session creation/enrollment.
- User registration/authentication.
- Membership Management.
- Active membership monitoring and report.
- Support of billing/charging model.

7.3.1 Activities and work items in ITU-T Q.8/17

ECTP (Enhanced Communications Transport Protocol)

- ECTP-1 (X.606): Reliable Multicast, Approved (2001.10).
- ECTP-2 (X.606.1): Multicast QoS Management (To be approved 2002).

GMP (Group Management Protocol)

- X.gmp: Session/membership Management.

RTM (Relayed Transport for Multicast)

- X.rtm: Hybrid delivery of broadband multicast traffic, based on unicast and multicast transports.

ECTP features interaction with IETF, and adopt the IETF TRACK approach for error control. ECTP distinctively provides:

- Tight Connection Control (ECTP-1)
  - Sender is at the heart of the multicast communications.
  - Session Creation/termination/pause/resume.
  - Connection Join via Sender (after authentication).
- QoS-based Session Management (ECTP-2)
  - Support of resource reservation (RSVP).
  - Session Monitoring and Maintenance.
  - Rate adaptation based on the monitored QoS status.

The ITU-T writes that basic multicast technologies are stable, like multicast routing, but there is still need to improve QoS management, Group management, etc. so as to accelerate multicast deployment.

There is also an identified need to develop an alternate solution for multicast, which ITU calls Relayed Transport Multicast (RTM), that based both on unicast and multicast.

Future works include development of multicast technologies and multicast control at transport layer:

- ECTP (Enhanced Multicast Transport Protocol).
- GMP (Group Management Protocol).
- Alternative delivery scheme for multicast.
- RTM (Relayed Transport for Multicast).

- Deployment of multicast services.
- Consolidation of business model.
- Broadband multimedia services.
- Multicast delivery services.
- Interaction between service and technologies.

References cited include ITU-T Recommendations:

- X.601: multi-peer communications.
- X.605: ECTS (Enhanced Communications Transport Services).
- X.606: ECTP-1 (part1).
- X.606.1: ECTP-2 (part 2).
- X.gmp: working document in Q.8/17.
- X.rtm: working document in Q.8/17.
- http://ectp.etri.re.kr.
- IETF WG.
- RMT WG.
- MBoneD WG.

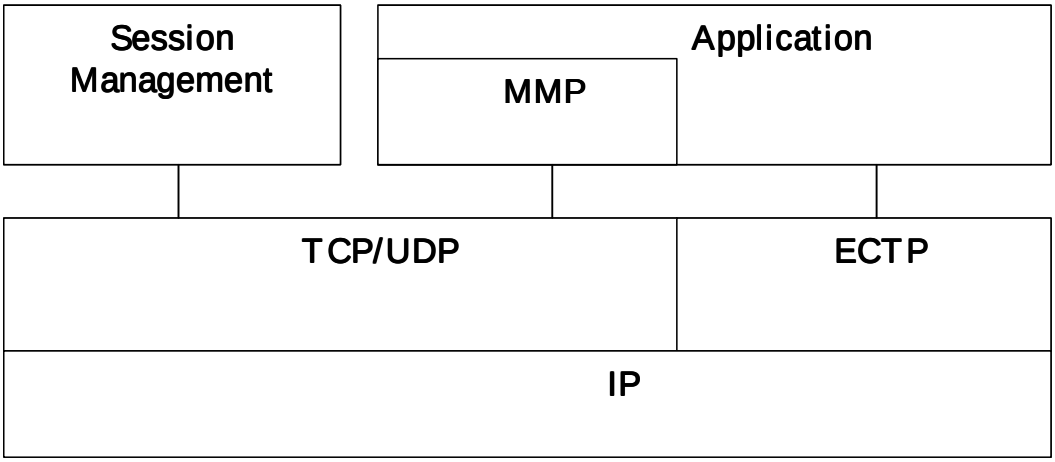

Figure 30: ITU group management overview

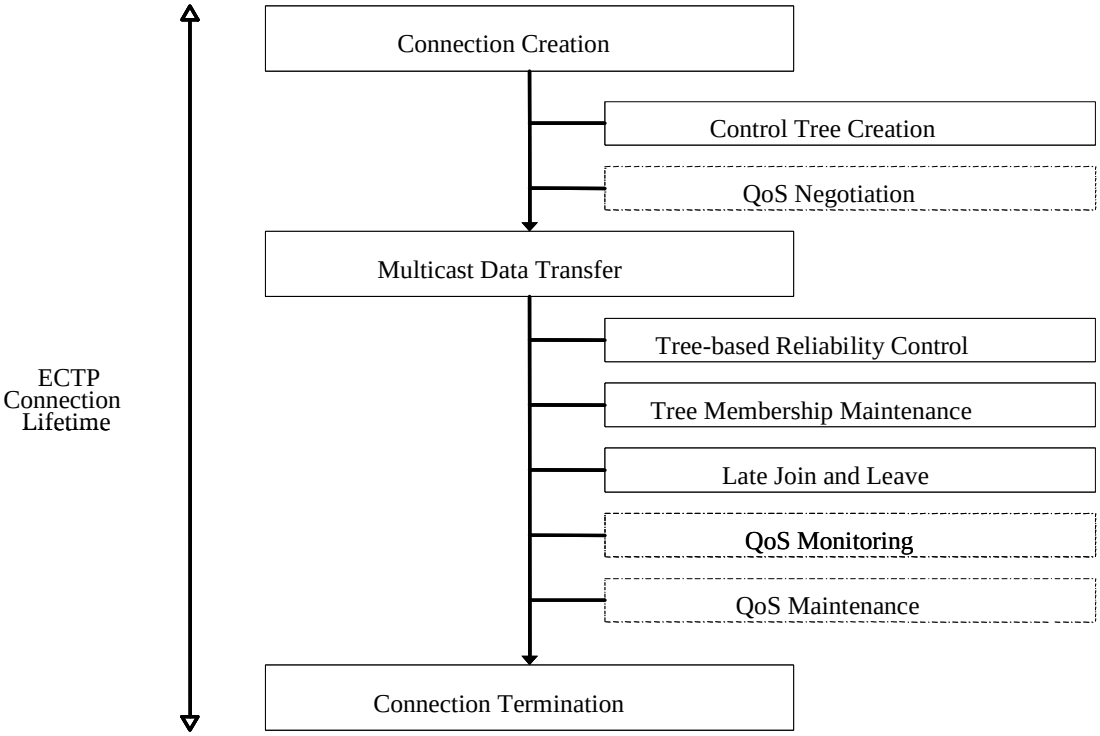

Figure 31: ITU ECTP QoS model

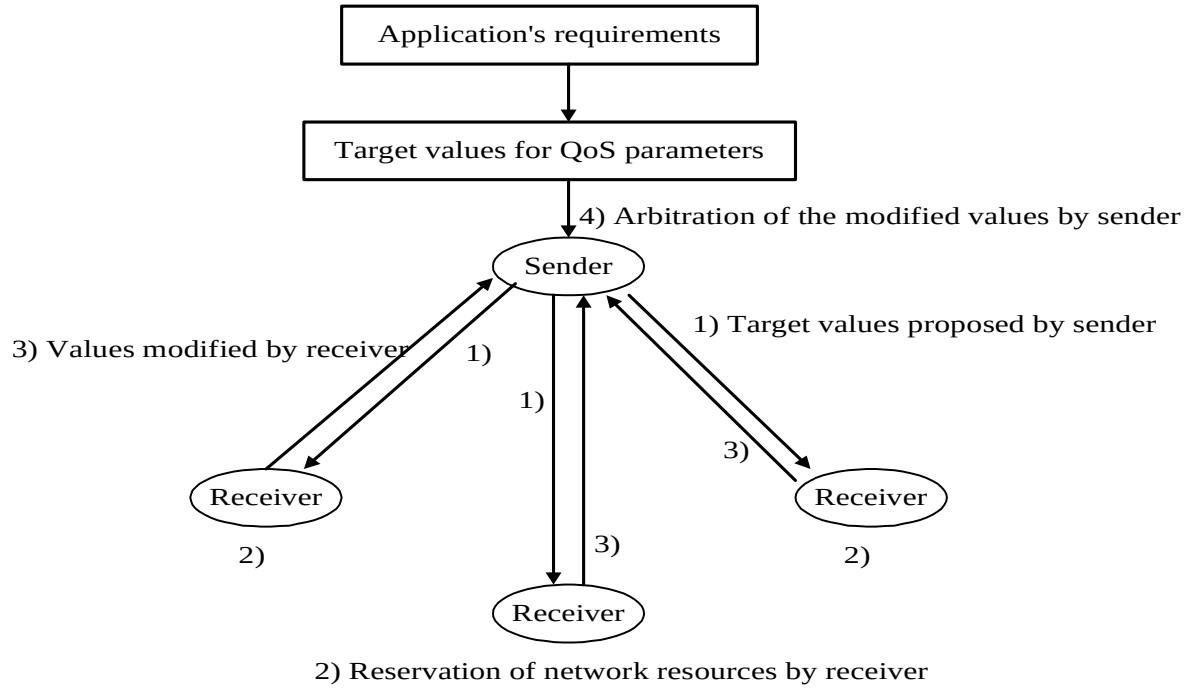

Figure 32: QoS negotiation policies

7.3.1.1 ECTS definition

ITU-T Recommendation X.605 ISO/IEC 13252: "INFORMATION TECHNOLOGY - ENHANCED COMMUNICATIONS TRANSPORT SERVICE DEFINITION".

This Recommendation defines an enhanced transport services for the next-generation Internet, named ECTS (Enhanced Communications Transport Service), which provides for a multicast data transfer capability and enhanced quality of service (QoS). X.ECTS defines a wide range of services ranging from unreliable unicast with best-effort QoS to reliable multicast with guaranteed QoS. In this way, X.ECTS is meant to provide for a uniform and universal service interface between transport protocols and applications of the present and the future information age, especially for those applications requiring versatile and powerful multimedia group communication capabilities underneath.

X.ects was approved at SG 7 plenary meeting in September 1998 as an ITU-T Recommendation. This was also approved as an ISO/IEC, in January 1999.

7.3.1.2 Multi-peer framework

ITU-T Draft Recommendation X.601: " MULTI-PEER COMMUNICATIONS FRAMEWORK.

This Recommendation provides the basic framework to specify services and protocols for multi-peer communications. The scope of this Recommendation is to define the basic concepts of group and various aspects of group communication, which are needed to specify specific services protocols for multi-peer communications.

X.601 was approved at SG 7 plenary meeting in March 2000 as an ITU-T Recommendation.

7.3.1.3 Simplex multicast transport

ITU-T Draft Recommendation X.606 | ISO/IEC 14476-1: "INFORMATION TECHNOLOGY - ENHANCED COMMUNICATIONS TRANSPORT PROTOCOL: SPECIFICATION of SIMPLEX MULTICAST TRANSPORT".

This Recommendation | International Standard specifies the Enhanced Communications Transport Protocol (ECTP), which is a transport protocol designed to support Internet multicast applications running over multicast-capable networks. ECTP operates over IPv4/IPv6 networks that have the IP multicast forwarding capability with the help of IGMP and IP multicast routing protocols. ECTP could possibly be provisioned over UDP. ECTP is targeted to support tightly controlled multicast connections. This first part of ECTP defines the protocol which provides reliability control in the simplex multicast case, adopting a tree-based approach.

QoS management functions for the simplex case will be defined in part 2 of the ECTP specification. Further parts of ECTP will be define reliability control and corresponding QoS management functions for the duplex case (parts 3 and 4) and the N-pled case (parts 5 and 6).

X.606 was approved by ITU-T SG 7 in 2001/10/29. This was also approved by JTC1 in 2002/01/15.

7.3.1.4 QoS for simplex multicast transport

ITU-T Draft Recommendation X.606.1 | ISO/IEC 14476-2: "INFORMATION TECHNOLOGY - ENHANCED COMMUNICATIONS TRANSPORT PROTOCOL: SPECIFICATION of QoS MANAGEMENT FOR SIMPLEX MULTICAST TRANSPORT".

This second part of ECTP defines the QoS management functions for the simplex multicast case. Further parts of ECTP will define reliability control and corresponding QoS management functions for the duplex case (parts 3 and 4) and the N-plex case (parts 5 and 6).

This specification describes the following QoS management operations, which are designed to establish and maintain a desirable quality of service: QoS negotiation, QoS monitoring, and QoS maintenance.

7.3.2 SG 13: Multi-protocol and IP-based networks and their internetworking

SG13 is responsible for studies relating to interworking of heterogeneous networks encompassing multiple domains, multiple protocols and innovative technologies with a goal to deliver high-quality, reliable networking.

Responsible for studies relating to:

- internetworking of heterogeneous networks encompassing multiple domains, multiple protocols, and innovative technologies;
- delivery high-quality, reliable networking;
- architecture, interworking and adaptation, end-to-end considerations, routing and requirements for transport.

Lead Study Group for:

- IP related matters.
- B-ISDN.
- Global Information Infrastructure.
- Satellite matters.

As SG13 is also Lead Study Group for IP related studies. SG13 has developed an ITU-T IP Project with the objective to cover all ITU activities on IP standardization.

The IP Project has a strong link with IETF:

- To avoid duplication of work and divergent standards.
- To collaborate where appropriate.

The IP Project is divided in 13 work areas:

- Area 1 - Integrated architecture.
- Area 2 - Impact to telecommunications access infrastructures of access to IP applications.
- Area 3 - Interworking between IP based network and switched-circuit networks, including wireless-based networks.
- Area 4 - Multimedia applications over IP.
- Area 5 - Numbering and addressing.
- Area 6 - Transport for IP-structured signals.
- Area 7 - Signalling support, IN and routing for services on IP-based networks.
- Area 8 - Performance.
- Area 9 - Integrated management of telecom and IP-based networks.
- Area 10 - Security aspects.
- Area 11 - Network capabilities including requirements for resource management.
- Area 12 - Operations And Maintenance (OAM) for IP.
- Area 13 - Utilization of IP v6 in telecommunication networks.

There is no area dedicated to satellites in the IP Project, but satellite issues are addressed in specific areas:

- Area 1: architectural aspects (application of satellites in the evolving network environments).
- Area 2: issues related to access to IP applications via satellite.
- Area 6: efficient ATM multicasting on satellite.
- Area 13: IPv6 satellite access network.

7.3.2.1 Next Generation Networks (NGN)

SG13 decided in Feb. 2002 to start the preparation of a new ITU Project on NGN, with the objective to respond to the demand from the market for standards, on a worldwide basis. NGN (Next Generation Network) is a concept widely used by network designers to describe future networks which should cope with the emerging situation in telecommunications:

- Open competition, total deregulation of markets.
- Explosion of data traffic due to the general use of Internet.
- User demand for new multimedia services and for mobility.

The Project should cover all ITU activities on NGN standardization, with active collaboration of involved SGs. The target date for first set of Recommendations on NGN is mid-2004 (end of study period). The view of SG13 that the NGN should be seen as the concrete realization of GII (Global Information Infrastructure) concepts which are defined in Recommendations from the Y series.

A list of capabilities of NGN is provided below:

- Creation, deployment and management of all kinds of services using all kinds of media.
- Architecture reflecting a clear decoupling between service functions and transport.
- Functional entities controlling policy, sessions, media, resources, service delivery and security to be distributed over the infrastructure.
- Interworking between NGN and existing networks to be based on gateways.
- Support of existing and NGN aware terminals.
- Migration of voice services with QoS and security.
- Generalized mobility (users and terminals).

In the draft version of the NGN 2004 Project description document, satellite aspects are not clearly addressed, because the description of study areas is still very general at this stage.

7.3.2.2 Study Group 13 Question 13 - Interoperability of Satellite and Terrestrial Networks

1) Background and justification

The explosive growth of data traffic has been driving changes in satellite communications. Satellite systems have been and are being planned, designed, and developed to offer services such as integrated broadband voice/data/video communications, mobile communications, and high speed internet access. Satellite systems also play an indispensable role in the convergence of telecommunication and broadcasting businesses. One major challenge facing the satellite industry today is to ensure interoperability of satellite and terrestrial telecommunications networks.

This question will focus on issues that affect interoperability of satellite and terrestrial packet-based networks.

2) Items for study

a) Develop architecture and reference models that integrate satellite systems with terrestrial telecommunications networks.

b) Develop new Recommendations in the areas of interconnecting packet-based networks (e.g. IP, B-ISDN, ISDN, and Frame Relay, etc.) using satellite.

c) Assess the operation of protocols (e.g. TCP/IP) when satellite is used as the transmission medium, and provide inputs to relevant Questions and Study Groups.

d) Study issues related to the use of satellite medium in an integrated telecommunications and broadcasting environment.

e) Identify existing ITU-T Recommendations that require improvement and revision in order to better interoperate with satellite systems. This effort will focus on ensuring established performance objectives do not preclude satellite systems.

3) Specific tasks and deadlines

- Revision of ITU-T Recommendation I.572 [33].
- Revision of ITU-T Recommendation I.571 [32].
- Specify network capabilities for interworking IP networks using satellites.
- Specify network capabilities and functional arrangements of heterogeneous networks for Tele-Broadcasting environments.

4) Relationships

- ITU-T Study Groups on general service aspects.
- ITU-T Study Groups on TMN.
- ITU-T Study Groups on data communications.
- ITU-T Study Groups on signalling aspects.
- ITU-T Study Groups on transmission performance.
- ITU-T Study Groups on multimedia services.
- ITU-R Study Groups on satellite aspects.
- IETF and ATM Forum on satellite issues.
- IETF and ATM Forum on satellite issues.

8 BSM multicast discussion

8.1 Gap analysis

The IETF has done and is doing a considerable amount of work on multicasting, and is likely to be the main standards body where multicast standards are developed also in the future.

The ITU is doing work related to multicast, and claims are that more work is needed on the transport protocol layer, i.e. layer 2.

There are no main findings of multicast work in DVB, but then again the basic DVB work is on layer 1, like the definition of DVB-S and DVB-RCS, although other layers are also discussed.

There seem to be no other body doing substantial work related specifically to satellites used in multicast scenarios, yet there has been identified a need for more work in the layer 2 domain.

We do not know the short or long term developments of multicast services. Multicast deployment has been delayed compared to expectations a few years ago. There are technical aspects that concern scalability and traffic control and with respect to services there may be a lack of multicast killer applications. There is also uncertainty of multicast business model. However, there are increasing needs for multicast deployment as part of broadband multimedia service such as web casting of broadband streaming media, remote education, stock-tickers, etc. This is a relatively clear business model.

- Deployment Scenarios for broadband multimedia in the short term indicate the continued use of unicast technologies with local multicast only in private networks, like potentially satellite networks.
- In the medium term one can see hybrid delivery of unicast and multicast with seamless transport via various technologies and perhaps tunneling or unicast over non-multicast regions. Satellite networks can locally also provide along with publicly available multicast content.
- In the longer term there will probably be full multicast over multicast-enabled regions and a significant expansion of multicast-enabled networks. In this term there will also evolve global markets for providing hi quality broadband multicast content over various multicast technologies.

Presently most multicast content is based on best effort delivery and the use of UDP. This is likely to change in the long run, and QoS based schemes will emerge. As various wireless technologies have become or are becoming more popular there is likely to be a set of dedicated transport protocols making efficient use of the radio spectrum whether it is satellite, UMTS, GSM or WLAN spectrum. These technologies will likely not be developed by ITU or IETF in general, with possible exceptions for the major technologies like 2G/3G technologies. It is hard to take general and specific satellite issues into account for non-satellite oriented bodies, like ITU and IETF, although it is likely that the satellite community will want to as far as possible be able to reuse the standards developed for the general telecommunications field to the extent possible.

ETSI should focus work on multicast on layer 2, but taking advantage of the SI-SAP definition and also taking into account the ability to provide local content for the satellite network alongside Internet multicast. Group management must also take into account at least the two basic groups business/professional and SOHO/consumer. In the latter case the satellite network is also likely to be the last mile network for the users.

Satellite specific transport layer functions (over BSM), QoS management (SI-SAP), security management (shared bearers) and group management (dynamic, internal/external groups) are areas where ETSI can and should play a role for the satellite community.

8.2 BSM multicast key findings

Subnetworks using shared channels (like satellite) are especially suitable for native multicasting, and should make every effort to support it. This involves designating a section of the subnetwork's own address space for multicasting. Satellite multicast specifications must scale gracefully with the number of users and not be limited by particular technology.

BSM multicast specifications will focus on the satellite independent layers, in consistency with the general BSM focus. Equipment manufactures and service providers need the capability to differentiate their products and services. Multicast concepts should also be "future-proof", as several different satellite technologies exist and will emerge, and that these need to be taken into account. ETSI satellite multicast specifications should take into account the work of the IETF, the ITU-T and ITU-R, TIA and DVB, and if possible and practical also input from prominent satellite multicast research projects.

Internet multicast applications over satellite will frequently interface with multicast groups where the majority of users (hosts) are outside of a specific satellite system, and need therefore to interoperate well. However, multicast groups limited within a satellite system are of interest in business applications and virtual private networks.

Disk and memory storage is relatively cheap and can be available in large amounts for gateways and in reasonable amounts for terminals. In this way caching both at network operator and customer premises can further increase the usefulness of multicast and should be included in a BSM multicast concept.

Security and key distribution are important issues in satellite multicast systems with shared bearers, and need specific solutions. Also the privacy of the end user needs to be protected by the satellite multicast protocols and any proxy use, i.e. the subscription to a specific multicast group should in general not be visible to the service provider.

The key findings in this report indicate that BSM multicast must interact externally with commonly used multicast protocols, specifically IGMP and PIM-SM. Interworking to others and also future protocols must be handled.

Flooding protocols for group management and routing are not desired, but rather user initiated versions that minimize the use of satellite spectrum for management. Several protocols have timers that regularly poll the status of host and routers, and these procedures can use significant amount of satellite network capacity if not properly handled.

A satellite multicast specification should aim to reach the receiver hosts with a minimum use of spectrum resources. This implies minimizing the cost associated in the delivery of the content; minimizing the overhead due to fragmentation, segmentation and transport of data and minimizing the spectrum resources required for management of multicast groups, i.e. key distribution etc. It should also minimize the delay for real time services over the satellite segment. A well defined specification should allow for future improvements in technology and protocols, and for future services that are not in use today. We also need to maximize the benefits a satellite multicast can contribute to multicasting, i.e. covering large regions, the ability to support high data rates.

When choosing or designing protocols and multicast solutions for satellite systems it is important to consider the following issues:

- Minimizing or avoiding control messages via satellite.
  - This can be done by pruning tree branches with no subscribers and using a receiver-driven tree to explicitly join and leave groups. It is also important to carefully choose suitable refreshing periods. It is important to choose a hierarchical architecture to delegate management and dynamics.
- Minimizing potential repair requests and packet retransmissions at satellite link level.
  - If possible try and aggregate repair requests into groups.
  - Use NACK instead of positive acknowledgement.
  - Use error correction coding to minimize ARQ and retransmissions.
- Minimize packet duplication over satellite.
  - Multimedia content is bandwidth demanding and frequency resources are always limited as often expensive to utilize on a point-to-point level.
  - Multicast content should be replicated only the necessary beams and only once to those.
- Allow for satellite or terrestrial return channel, and even options for no return channel.
  - If possible and available terrestrial return links sometimes offer advantages.
- Allow for public domain Internet multicast groups, closed groups for the BSM networks and restricted groups for corporate user groups, and be able to provide high levels of security when required.

A BSM multicast concept should not require any particular action from senders outside the satellite system to take satellite specific issues into account.

Emerging standards should permit the possibility of making a simpler BSM network with a limited edge support at a simpler terminal, e.g. only supporting IGMP version 2. The solution within BSM network depends on the requirement of the edge protocols. Such a solution may be as simple as implementing stub multicast routing in the terminal by introducing an IGMP version 2 proxy with special constraints for BSM network signalling load reduction. This solution may effectively serve many applications used in a scenario with satellite terminals acting only as multicast network stubs.

The BSM multicast concept we are looking at below is an "idealized" concept, which can support:

- hosts as both sources and destinations connected to a satellite terminal;
- subnetworks with PIM-SM routers for larger networks and in the gateways for communication with external hosts (on the public internet).

In the below scenario the satellite would replicate messages as required, according to routing trees, and perform group management as required as well. The reason for calling it "ideal" is that when replication can be done in the satellite then the absolute minimum amount of replication takes place, preserving the maximum of spectrum. Replication takes place as close to the destination as possible. However, this spectrum-ideal approach may not be ideal in terms of cost and complexity for a practical system. In most practical systems, at least in the short term, full router functionality and group management functions will probably not be found on board the satellite(s). In spite of that the figure may give us something to strive towards when multicast specifications are being developed; and at the same time presenting a simplified diagram for our goals.

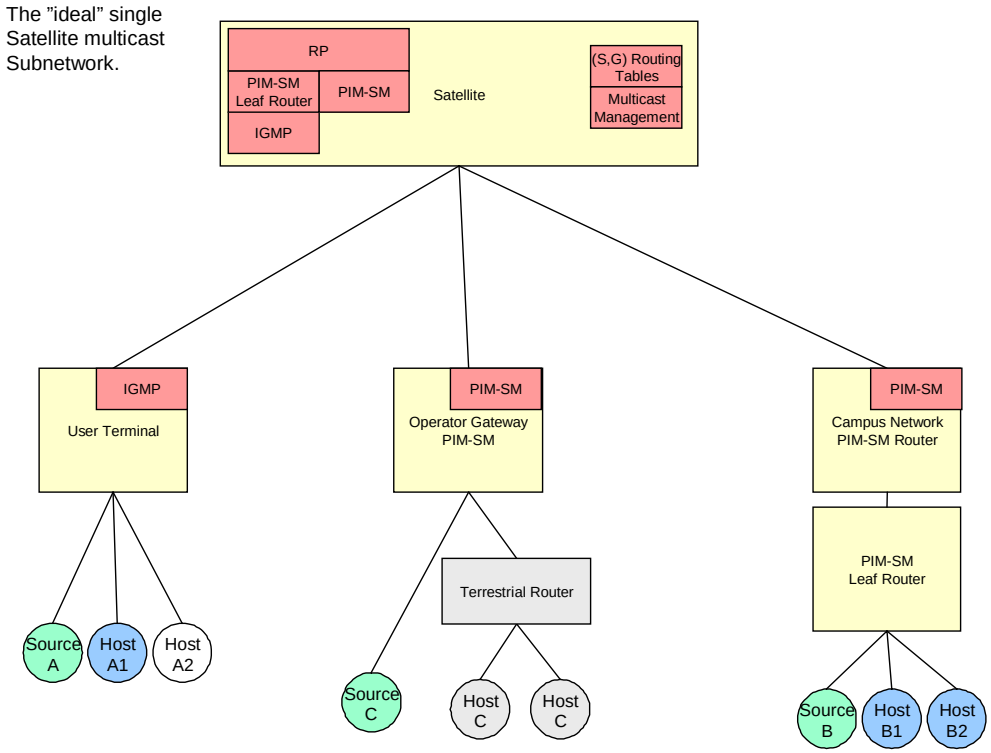

Figure 33: The "Ideal" BSM multicast concept, where the satellite segment connects to all destinations

The most relevant multicast protocols to use are IGMP and PIM-SM, as they are commonly found and used. Also required would be the SSM, source specific multicast.

Routing trees must be built depending on the BSM network architecture, but one should assume that every spot-beam is a branch in a routing tree. Any PIM-SM router in a beam would not see other routers, and as such any one would be a designed router. Information on available multicast groups must be available for the routers either locally (e.g. via a broadcast) or at the next branching point which would be the satellite or the gateway depending on the OBP capabilities of the satellites.

Address space and mapping to layer 2 addresses will differ, and there are several options. In a system with limited addressing capabilities the solution is to let terminals decode the data and filter at the IP level. IP multicast address mapping is essential in satellite multicasting.

Further issues identified are:

- Specifications should primarily be based on a relevant set of IETF RFCs (or upcoming RFCs) and use these consistently, with added "inter-layer" assistance from ETSI specifications.
- Computer hosts connected to the satellite terminal should subscribe to multicast groups using standard IP protocols and software at application layer.
- Satellite gateways and service providers should be able to provide additional (value added) multicast content to the satellite network subscribers.

- The satellite network itself can decide the underlying link layer protocols and potential use of simplex or duplex communication. It can e.g. be one of several BSM alternatives including also DVB-RCS.
- Hosts on the satellite network can also be the originator of multicast messages.
- Specifications should be able to provide IP-sec or similar end to end security.

Things ETSI should not do include:

- Defining "competing" multicast protocols with IETF standards.
- Defining standards locked to specific air interfaces, possibly with the DVB exception.
- Defining standards that are out of line with technology trends and business models.

8.3 Source group table management

For every source there will a group that is associated with it. For every group there will be a routing table. The table will need to track the following entries in the general case (for some systems not all parameters may be relevant), assuming the multicast source enters the BSM network at some point and is destined for a set of terminals and their hosts:

- Gateway: Which gateway, and how to reach it. This includes the routing to the gateway in question which can be basically anywhere on earth, and where the right gateway may be reached via satellites or terrestrial networks. In an OBP system with on-board switching/routing capabilities and inter-satellite links, any one gateway may perhaps be used. However, if there are several options, somehow a decision needs to be made. For systems without inter-satellite links a specific gateway choice may need to be made.
- Satellite: Which satellite and how to reach it. Constellations can vary from single GEO to LEO with tens of satellites.
- Spotbeam: The right spotbeam needs to be decided. Future satellites will have hundred of spot beams.
- Carrier frequency with in the spotbeam must be selected. Every spotbeam may have more than one carrier, or even no permanent carriers for scanning spot beams.
- Terminal timeslot: the timeslot(s) on the chosen carrier must in some cases be decided. This will depend on if the terminals actually decode everything they receive or simply listen on specific time instances. Both categories of systems exist. (For a CDMA system the spreading code would have to be decided).
- Host or next routers: The multicast content may need to be routed to one of a few hosts on a local premises network, or to another router on a campus network.

If the originator of the message is a source connected to a satellite terminal, then the message may have to be tunneled to a starting point (in PIM-SM in would be the rendezvous point) via the gateway the terminal communicates with. If the satellite is OBP it could simply forward the data to the satellite which in turn would multicast it to the terminals it could se, and then route it to other satellites and gateways to reach the remaining set of receivers (which could be both on the BSM net and external on other satellite networks or terrestrial networks).

Management of the entries in a routing table will generally be via normal join and leave messages from hosts. Such messages can for instance be generated as a response to a customer action. Joining and leaving groups may be via a web page interface. However for terminals without a return channel, static or semi-static entries may be entered by an operator.

To conserve bandwidth on the satellite network, there may be little or no need to regularly ask terminals if they want to join specific groups.

8.4 Address space management

It is likely that multicast replication will be done at layer 2, and transport will be based on L2 addresses. If for instance all multicast groups shared the same PID or L2 address, then all groups would be replicated even if they were not requested, and filtering would be done in the receivers. However, this would be a way to waste satellite spectrum resources, and more efficient ways should be our target.

A BSM system must decide how L3multicast addresses are mapped to L2 multicast addresses. One BSM system may have a larger L2 address space than another, but there would in any case be a need for address resolution between the Class D multicast IP address and the L2 multicast group ID. There are 28 bits of unique address space for an IP multicast address (32 bits minus the first 4 bits containing the 1110 class D prefix.

For Ethernet only 23 bits used for mapping into the IEEE MAC address, thus 5 bits of overlap remain. This means that multiple layer 3 multicast addresses can map to the same layer 2 multicast address. Similar concepts need to be used for BSM multicast unless the L2 address space is larger than the IP multicast address space. The number of needed overlapping bits will however vary with the size of the L2 space, but a few guidelines are still possible to decide in a specification. However, the presence of non-unique L2 multicast addresses prevents IP multicast capable switches from building independent L2 layer multicast delivery trees for each IP multicast group thus causing some overlapping delivery of unnecessary IP multicast packets. To maximize L2 multicast efficiency, delivery of IP multicast groups must be based on a unique mapping to L2 MAC address. In a packet switched system, there will have to be a trade-off as to how much addressing overhead and overhead that should be carried with each packet compared to the actual payload size.

To support unique L3 to L2 mapping efficiently, large packets are favoured to reduce the relative overhead. Addresses should be designed in a way that allows efficient processing in terminals. A BSM multicast system requires efficient L3-L2 mappings.

8.5 Brief BSM Multicast Proxy description

A BSM Multicast Proxy (BSM-MP) should take advantage of the specific satellite systems strong features at the various layers, yet making different satellite system architectures basically invisible to the outside world. BSM-MP will know both the specifics of the underlying satellite system as well as the requirements of the services in question.

BSM-MP systems should be defined to work both on systems with satellite return channels, on systems with terrestrial return channel and on systems without return channels. If a source e.g. requires reliable multicast, then BSM-MP will take the capabilities of the underlying system into account when deciding the means of delivery (coding, modulation, possible ARQ, repeated transmissions, QoS, etc.).

BSM-MP will allow an underlying middleware function handling functions at the L2 level, which needs to be designed for each satellite system in question, to optimize the use of its resources, and offer a consistent set of capabilities and signalling of these capabilities to the outside world. BSM-MP will therefore work to interact with external multicast input and output with the whole satellite system as a sub-network.

In general, BSM-MP will thus act as a multicast enabled router seen from the core network side. External multicast sources will find the BSM-MP router, as any other multicast enabled router, and forward the multicast content to it. The BSM-MP will in turn deliver it to the terminals that shall have it.

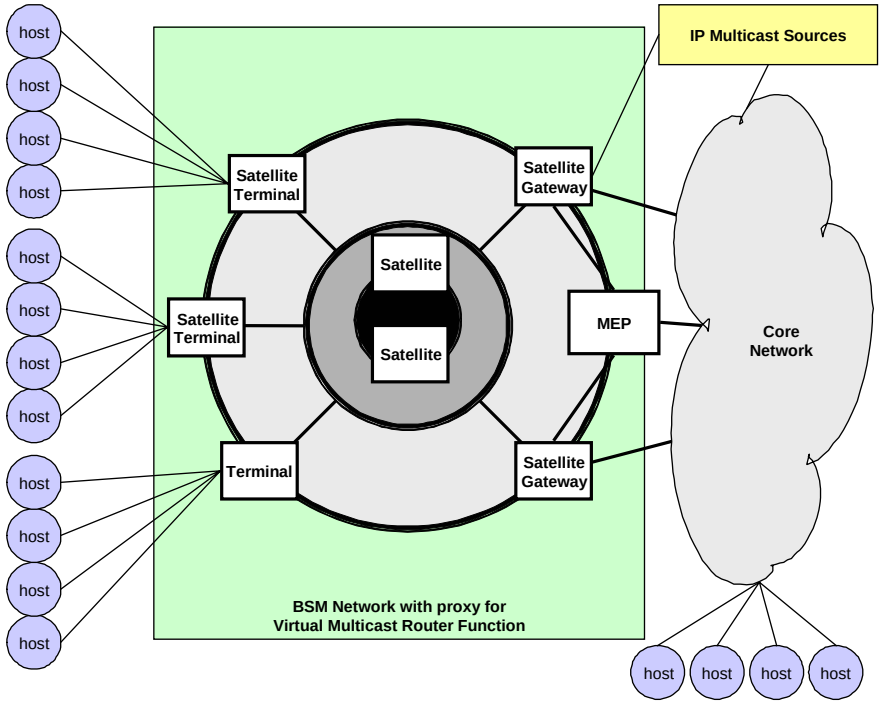

Figure 34: The BSM-MP (Multicast Proxy) functionality will appear as "just another router" and a sub-network seen from the core network. Sources can be both external and internal.

The satellite gateways or Network Control Centers (NCC) (depending on the architecture of the network) will be the router interfaces to external multicast groups on the Internet. Internal groups can originate also at the satellite terminals.

The BSM-MP will offer Internet multicast groups to the BSM sub-network, but there must also be the capability to filter groups (e.g. to reject spamming) and add group (e.g closed content).

The BSM-MP will act as a multicast manager. The required multicast manager functionality will be split between the different components as NCC (MEP/NEP), gateways and multi-user satellite terminals. For instance should terminals replicate messages locally when required, and be able to operate with the proposed protocols. A BSM-MP may take advantage of a matched set of L2/L3 functions at both the gateway and the terminals. This distinguishes it from traditional multicast standards, where routers and hosts basically are independent. In BSM system, the terminals will in any case be matched to the system in question (protocols, air interface, etc.), and BSM Multicast (BSM-M) adds on matched multicast functionality. BSM-M will in particular take advantage of the SI-SAP layer defined in BSM.

Terminals will normally subscribe via IGMP when the BSM network is an edge network. The BSM-MP will be seen as a PIM-SM router, but every spotbeam is generally seen as another subnetwork. With a PIM-SM router concept the logical satellite system Multicast Entry Point (MEP), for instance located in the NCC, would act as the Rendezvous Point (RP).

BSM-MP will control the replication processes, either directly or indirectly (like when content is replicated in satellites). It will also control the management and allocation of addresses and the session management.

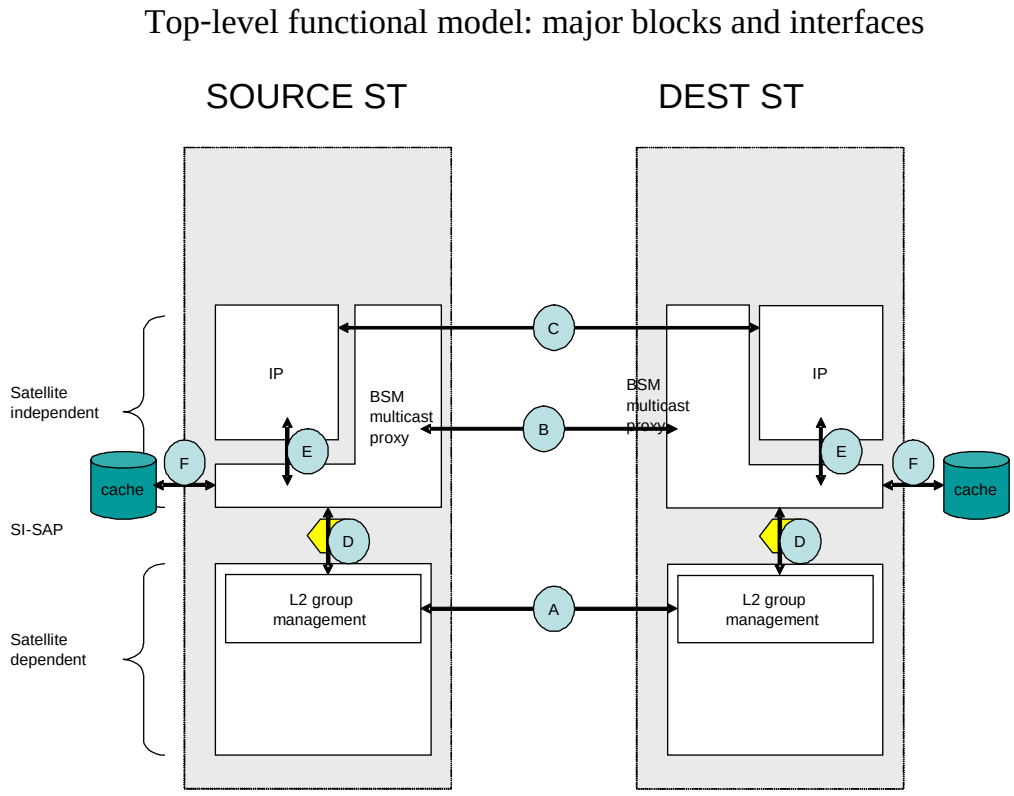

Figure 35: BSM proxy concept

The following interfaces are shown:

A: The peer to peer interface between the satellite dependent layer 2 entities.

B: The peer to peer interface between the proxy at the source and the proxy at the destination. This is a logical interface for proxy packets.

C: The peer to peer interface between the IP layer at the source and the IP layer at the destination. This is a logical interface for non-proxied IP packets.

D: The Satellite Independent Service Access Point (SI-SAP).

E: The interface between the IP layer and the proxy function. This is a logical interface, where proxied packets are forwarded to the local proxy function.

F: The logical interface between the proxy function and a local cache.

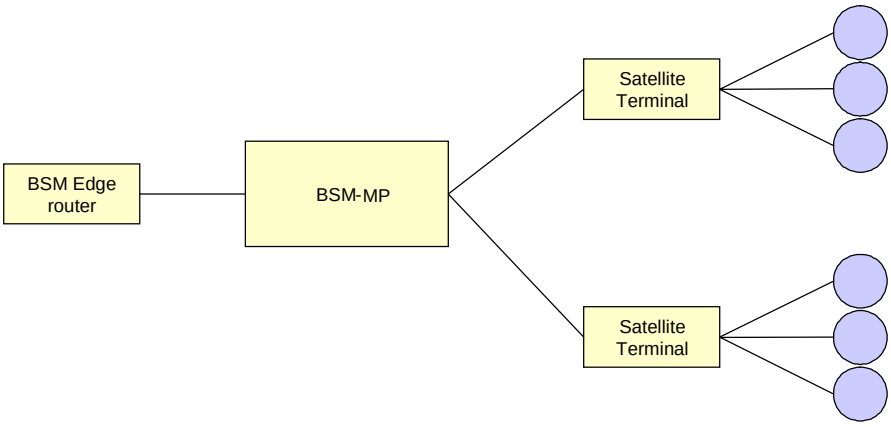

Figure 36: BSM-MP will in principle be a "black box" function between edge routers and terminal routers, that takes care of the technicalities of BSM multicasting for a given satellite system

Seen from the operator side, BSM-M will consider the capabilities of the underlying satellite network and forward massages to the correct nodes, i.e. gateways and terminals. Further, BSM-M can allow for a Multicast RRM function that carefully decides where messages shall be mapped to replicated, and also replicates the messages in question. BSM-M will not segment messages into layer 2 messages, but will in general depend on such a function being present.

Seen from the satellite terminal side, BSM-M will be a point for subscription of multicast groups. BSM-M will appear as a multicast router, and will allow mass market subscribers (i.e. consumers etc.) to receive (and send) multicast data over the satellite network with standard protocols (principally IGMP) on the host. The satellite terminal may function as a multicast router as well. A multicast manager function in the (gateway and terminal) can be used to optimize the use of spectrum resources. Subscribing to multicast groups will be a normal procedure with standard tools, and BSM-M will in this context be invisible.

For large scale business applications communication through corporate gateways (a matter of definition) BSM-M can allow the use of PIM-SM over the air, assuming the customer premises equipment has a multicast enabled router and a local network connected to it.

For a service provider, BSM-M will be a connection point for additional value added services, such as closed multicast groups for the satellite network in question (dedicated content, movies, concerts etc.). BSM-M will also handle local and closed multicast groups that for example can be set up by an organization with terminals at different sites. With these capabilities, BSM-M can also offer external multicast sources the capability to provide specific satellite versions of multicast content, tailored to the capabilities of the satellite network.

BSM-M will support reliable multicast, and selective repair request. The process of the repair request is yet to be defined.

BSM-M will also (in the future) be able to support QoS for multicast. Initially, such QoS support may be for the locally added sources (if any) as these would not depend on the capabilities of other multicast routers along the way from the source.

BSM-M will handle caching of non-real time multicast content when these capabilities are beneficial.

The added complexity in the terminals to support BSM-M should be kept to a minimum, specifically in terms of configuration and maintenance for the end users, but also in terms of additional HW and SW cost.

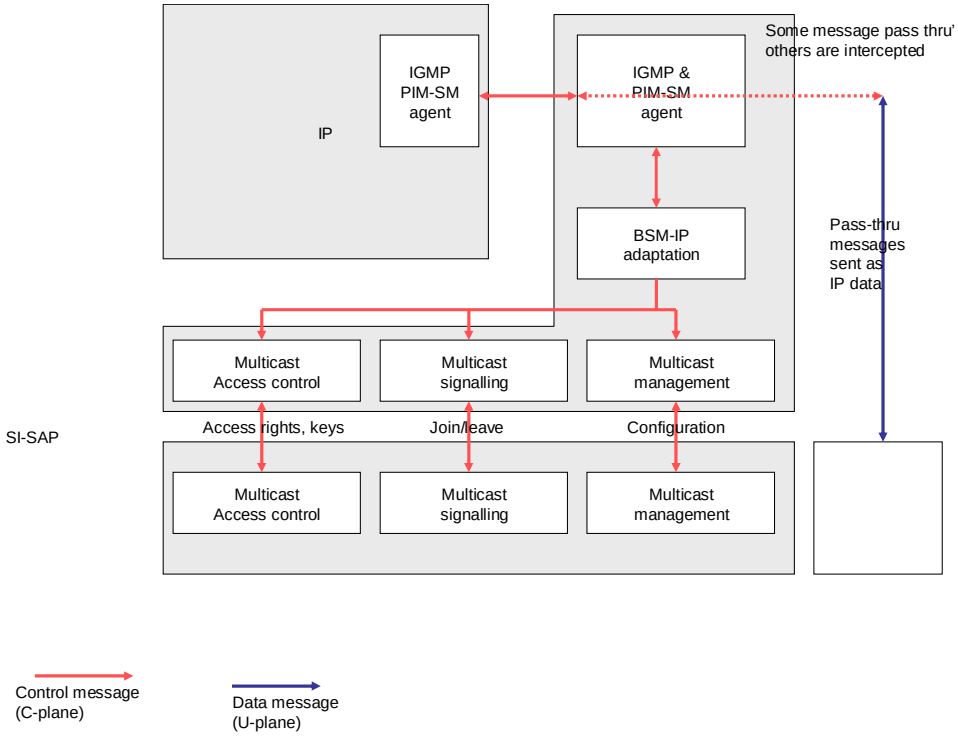

Figure 37: A brief look at the BSM Multicast proxy functionality

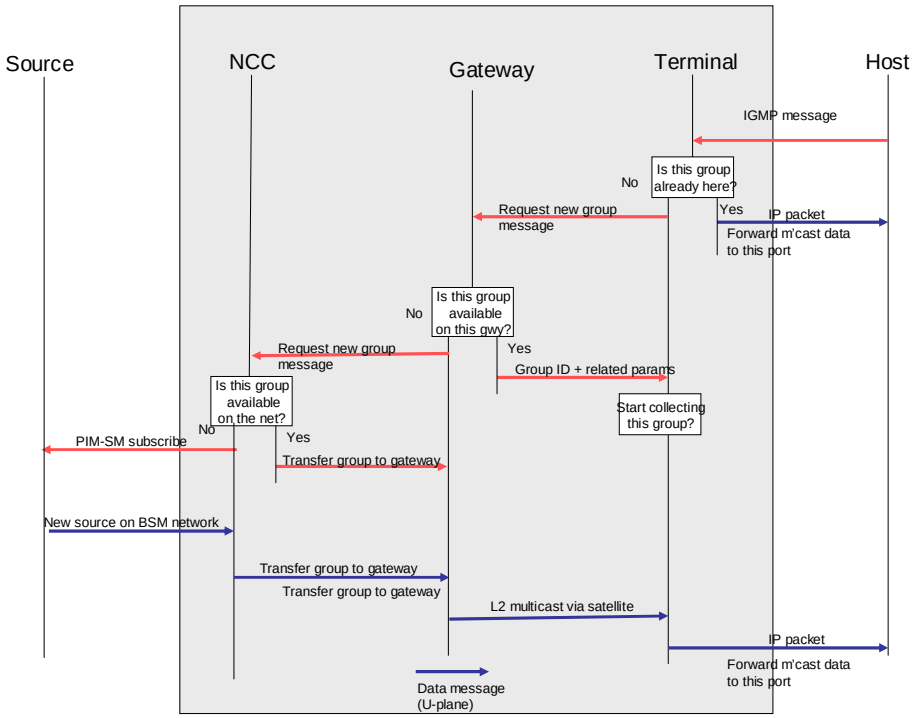

Figure 38: Interworking example

8.5.1 Basic needs and optional features

This clause will discuss BSM multicast basic needs, the desirable features, things that are optional, things that should not be done.

The basic needs consist of:

- The availability of a multicast source and knowledge of where it is connected and the intended "scope" i.e. where in the network it is intended to be delivered to. This may be dependent on the source address, or contributing network.
- A way to manage source/group memberships and addressing, i.e. a way to find out which down-links need to receive this multicast transmission and to address the receivers.
- A way to transfer the multicast data from the source to the destinations.
- A management plane for the multicast.

The source could be a publicly available source, in which case ETSI needs to do nothing specific. However, the source should also be possible to define locally to the BSM system in question, and in this case ETSI should define how such local sources should integrate with the other multicast content and at what level. Most probably it should be visible from the hosts connected to satellite terminals via commonly used multicast software, such as SDR, or one could use a web page, but this needs to be further looked into. It may also be possible that some dedicated version of the multicast software need run on the host.

Groups can be managed with common IETF protocols like IGMP and PIM-SM. However, group memberships at L2 must be subject to ETSI specifications. Thus the mapping tools, which can be simple (and perhaps not so efficient for complex satellite systems) or more advanced, depending on the requirements for spectral efficiency and other factors. Other factors include scalability, handling of dynamic groups and more.

Data transfer should be efficient, and that immediately distinguishes it from the plain ability to offer the service of transferring the data. Data transfer could for instance be unicast or broadcast, but none of these really take advantage of multicasting. Such techniques are not primarily what we are looking for in the medium to long term. In fact, the multicast data should only require capacity in those beams that have subscribers to the content, and further, preferably, also require a minimum number of uplinks, which is specifically relevant for OBP satellites. If the content requires reliable multicast as opposed to best effort, then the underlying satellite multicast system must decide issues like coding, modulation, which can differ in every beam, and possible the use of ARQ protocols.

The desirable features include:

- Optimum spectrum efficiency transport methods.
- QoS support.
- Support for IPsec or other security functions.
- Multiple source multicasting.
- Multicast content caching in gateways and/or terminals.
- The ability to manage the multicast process and place operational procedures to deny/allow sources to use specific uplink terminals, and possibly to control their transmit rate.

The key thing is to have a good efficiency, taking advantage of a high "N" in "1: N" multicast transmission, and thus sharing the cost of the use of the satellite resources by a large number of users. If N is large the major gain in efficiency may already be achieved. Optimality is nice, but not a firm requirement. Another example is that if there usually are a number of multicast recipients in every beam the perhaps the content just can be broadcast to all beams even if some capacity is wasted. Yet another example is that even if the satellite can replicate content and one thus can minimize the number of uplink copies this may add some complexity to the system which gives less payback than the basic initial efforts. However, the optimality is desired if it can come at a reasonable cost.

QoS support will probably be a firm requirement in the future, but at the moment multicast services are in general not supported with QoS. Therefore QoS is a desirable feature in the short to medium term.

The ability to provide secure data connections is crucial for some users. Initially in a system, multicast to secure users can potentially be handled by unicast IPsec given that the number of users is low. However, this is not a good solution in the medium or long term, and it is not at all a multicast solution. Toe support closed user groups and multicast on such groups it may be essential to provide a level of security above simple group membership management, thus some form of encryption, possibly specific to the BSM system in question, will at some point in time be required.

Caching can be used to provide several benefits for users, but also to use time as a factor for integrating the number of interested parties to a higher number, and thus provide lower cost for the transmission part. Caching can be either at the gateway(s), and can take the form of edge-casting, or to the individual terminals. If both options are available the choice of which storage to use would depend on the amount of storage required (it will generally be larger in the gateways), the type of content and the general interest for the content. Also, the ownership and the business case for delivering the content may influence the choice; i.e. it may depend on what the decisions are to the right of ownership to the content on the user terminals. Gateway caching may also be used if terminals are unavailable, for instance due to rain-fade, or in the case of nomadic (portable) terminals, they may not be logged on at the moment. "Recasting" the content may be attempted once later, depending on the number of parties that missed it.

Multiple source multicasting is useful in multi-party conferencing situations, for instance. However the resource management and the congestion control algorithms that would be required for bursty and non-predictable sources would presently seem prohibitive. However, can the problem be solved in a reasonable way, then there are applications that would be able to take advantage of such solutions.

Some protocols, such as IGMPv2, expect each terminal to generate multicast traffic. However, control traffic is usually expected to be light. One solution is to forward this to a gateway and then rebroadcast it to all other terminals. This is simple, but incurs a double satellite delay, and also requires measures to stop terminals receiving the packets they have sent!

8.5.2 BSM Multicast Networks

A multicast implementation relies upon a receiver-based protocol and multicast-aware/enabled switches and routers to replicate a single stream to multiple destinations. The architecture of the BSM multicast network is important, and it must also be seen in relating to the customer premises network. For a significant part of the multicast hosts the satellite network is also an edge network. This allows some more freedom in design. Once a multicast packet is delivered to the entry interface of the BSM network router, the BSM network handles its delivery to the receiver for the edge network segment.

Multicast is by nature a hierarchical protocol. Therefore, multicast will be easier to implement in hierarchical networks rather than flat, star, or bus networks. A BSM satellite network is a hierarchical network, as shown in figures 39 to 42. All relevant portions of the network should be multicast-aware, and then ensuring that each piece of equipment on those network legs is IGMP- and multicast-aware, is important. If each switch/router is IGMP-aware, then only the end stations that request a multicast stream will receive it. If they are not IGMP-aware, the switch/router will replicate the stream to every interface in order to ensure that it reaches its destination.

In addition to the architecture, a customer needs to determine which routing protocol they wish to use for multicast traffic. PIM (sparse-mode and dense-mode) are commonly chosen and BSM multicast protocols need to supports such network infrastructure. The customer will likely need a multicast-enabled router within the network to manage multicast routes, as well as to act as a rendezvous point for local traffic if PIM sparse-mode is used.

Figure 39 shows a multicast hierarchy, and figures 40 to 42 are modified to show reference architectures for BSM networks. Note that at this level there is no decision as to what level (L2/L3) the messages shall be redirected at. That can happen at routing level or at switching level.

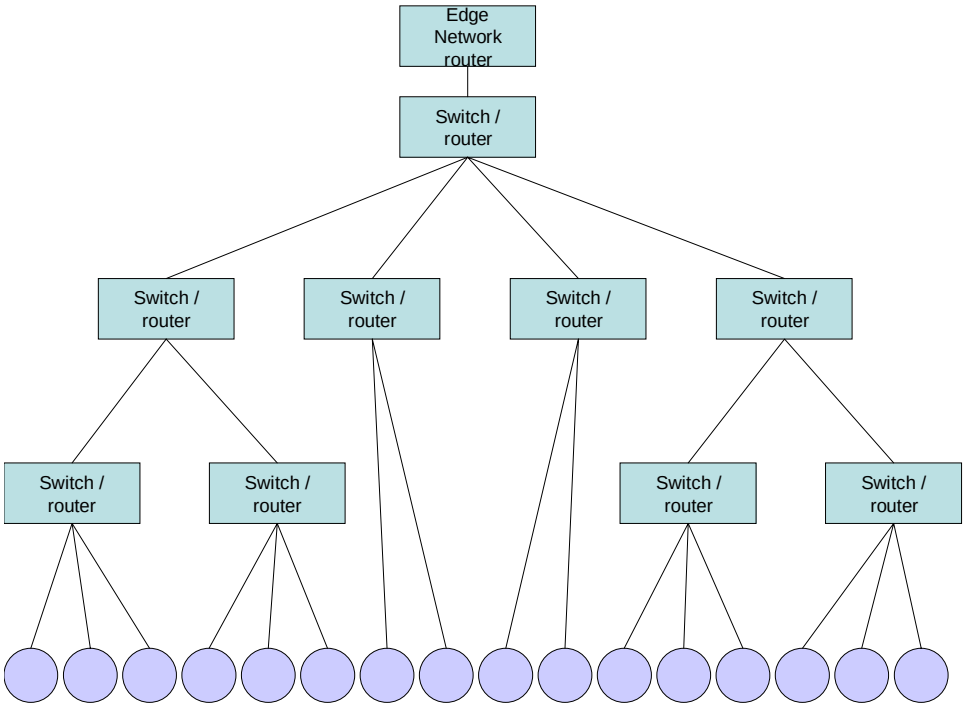

Figure 39: Multicast hierarchical tree architecture

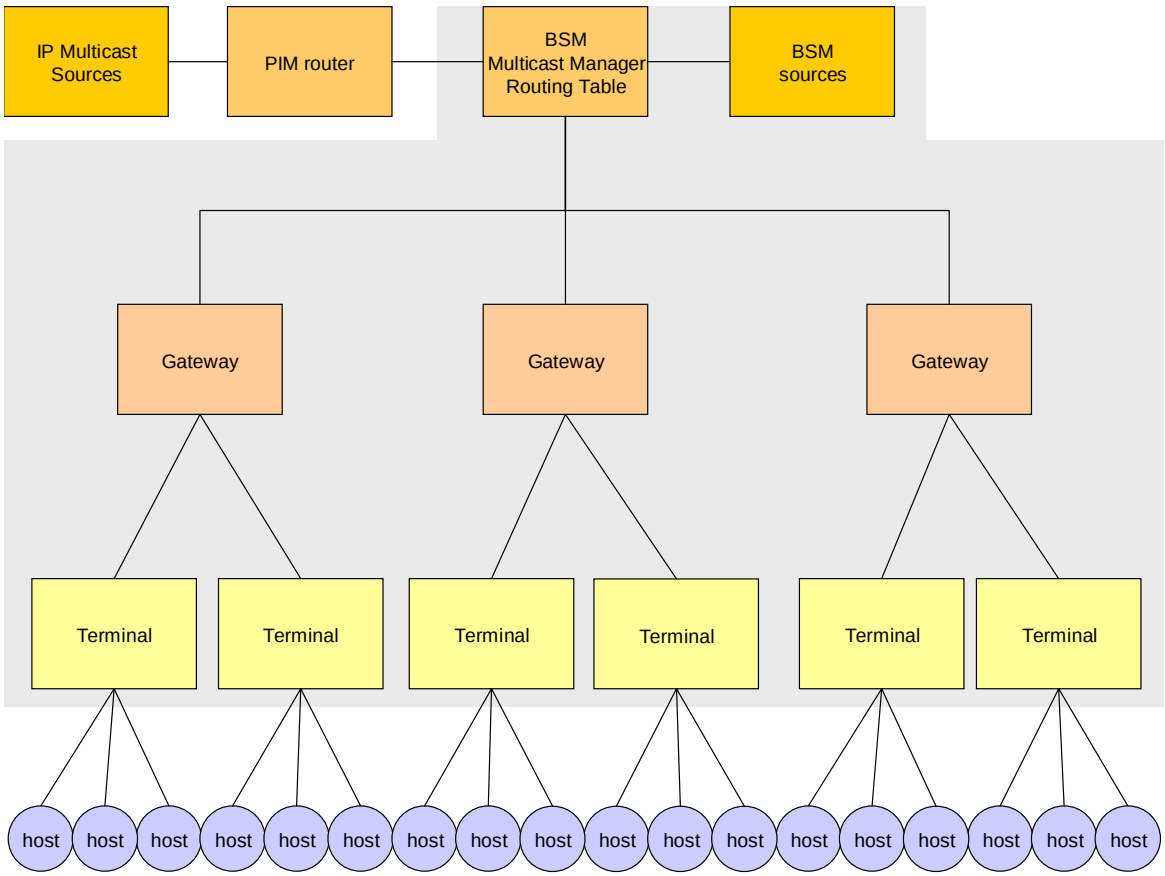

Figure 40: Multicast hierarchical reference architecture for a BSM network with bent-pipe satellite

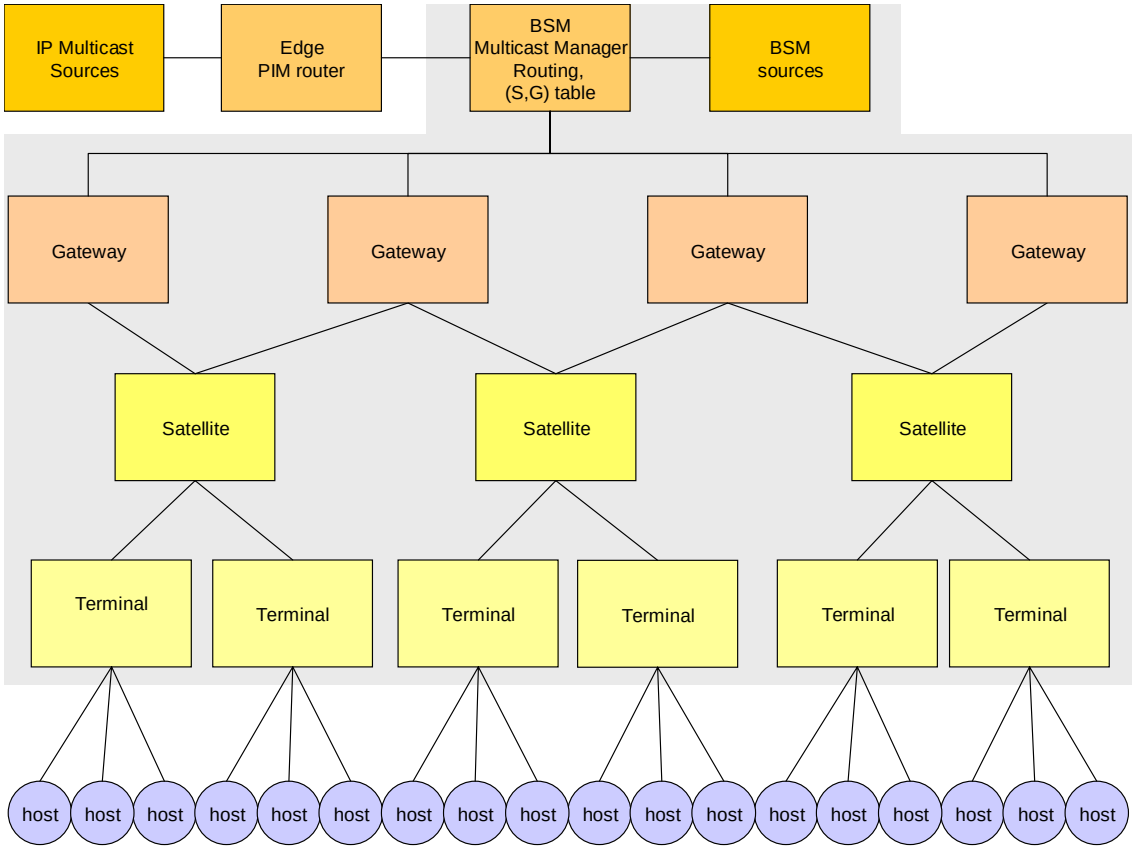

Figure 41: Multicast hierarchical reference architecture for a BSM network with regenerative satellite

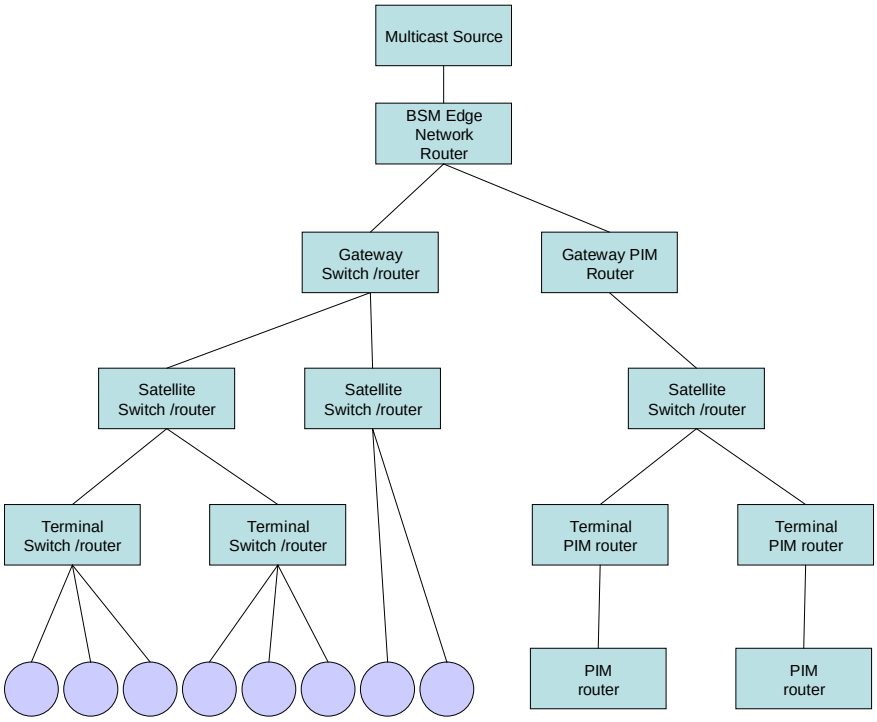

Figure 42: Multicast hierarchical architecture example for a BSM network with regenerative satellite and a combination of IGMP and PIM routing (or switching)

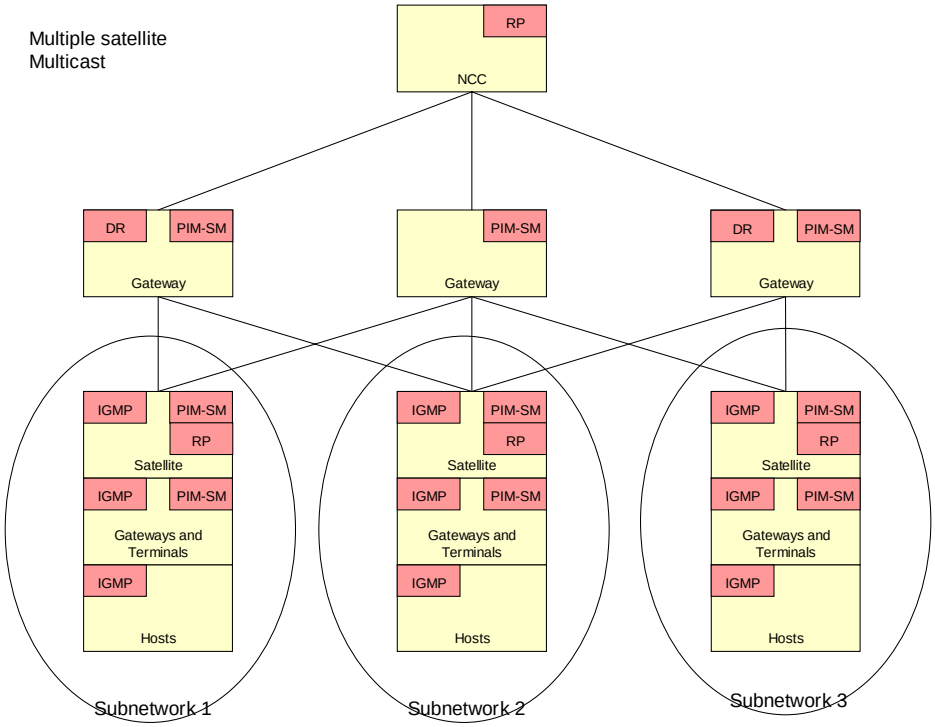

Figure 43: BSM multicast and sub-networks

The following network scenarios are required to be included as scenarios for satellite multicast standards:

- Open, as in the Internet - where there is no control over the entire network.
- Closed, as in an Intranet - where there can be control over total network and thus also the investments required for solutions and the efforts required running the network.

In open networks, interworking with other standards is crucial, and in general, given the number of hosts on the Internet, the origin of the multicast will not be within the satellite system in question. Rather, the users on that system in general subscribe to different groups. However, the different satellite hosts can be the origin of multicasting, but in this case, the general assumption will be that the recipients are not on the satellite network. Again, there is also a good chance that some recipients will be on the satellite network.

For closed networks, it is more likely that the majority of the receiver is on the same system, i.e. the same satellite system. Also, specific services could be envisioned, that are not IP specific. Applications include such as corporate web casts, video conferencing, distance learning etc.

- Return path (via satellite, via terrestrial or non-existent. Networks are fundamentally different depending on if they have a return path or not).
- Availability (depends on fading, among other things. May require dynamic routing).
- Transfer capacity (dynamic routing may increase the average network load).
- Delay and jitter (influenced by terrestrial routing, inter-satellite links and multiple satellite hops).
- Gateway selection in multiple gateways (part of routing).
- Network assistance, the use of different layers.

There can be a combination of the two types, where satellite networks offer access to public domain multicast groups an in addition they offer value added content or services for the users on the satellite segment.

8.5.3 BSM multicast protocols

It is proposed to use an interworking function at the IGMP level for consumer terminals. A draft version of the protocol stack is illustrated below.

The use of multicast proxies in a satellite network implies that multicast messages are intercepted, and handled by one or more multicast proxies in the satellite network. The use of multicast proxies will offer compliance with normal Internet multicast protocols yet allow special optimization or the space segment. Multicast proxies are used also in the satellite terminals to interface with hosts connected to the satellite terminal.

Key functions that are to be handled include L3-L2 addressing and mapping, replication, source-group routing tree construction and IGMP and PIM-SM timer spoofing.

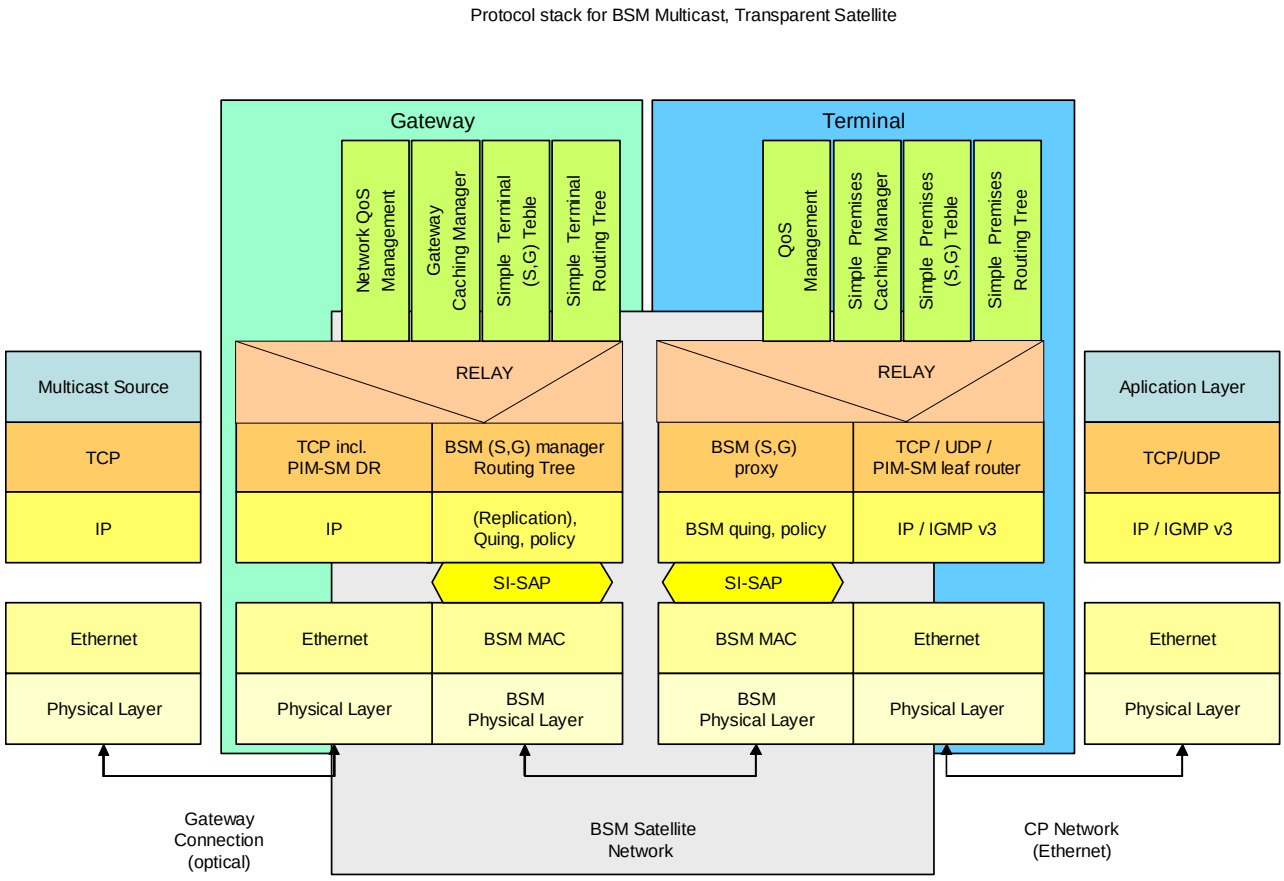

Figure 44: Draft sketch of a BSM Multicast protocol stack concept for bent pipe satellite

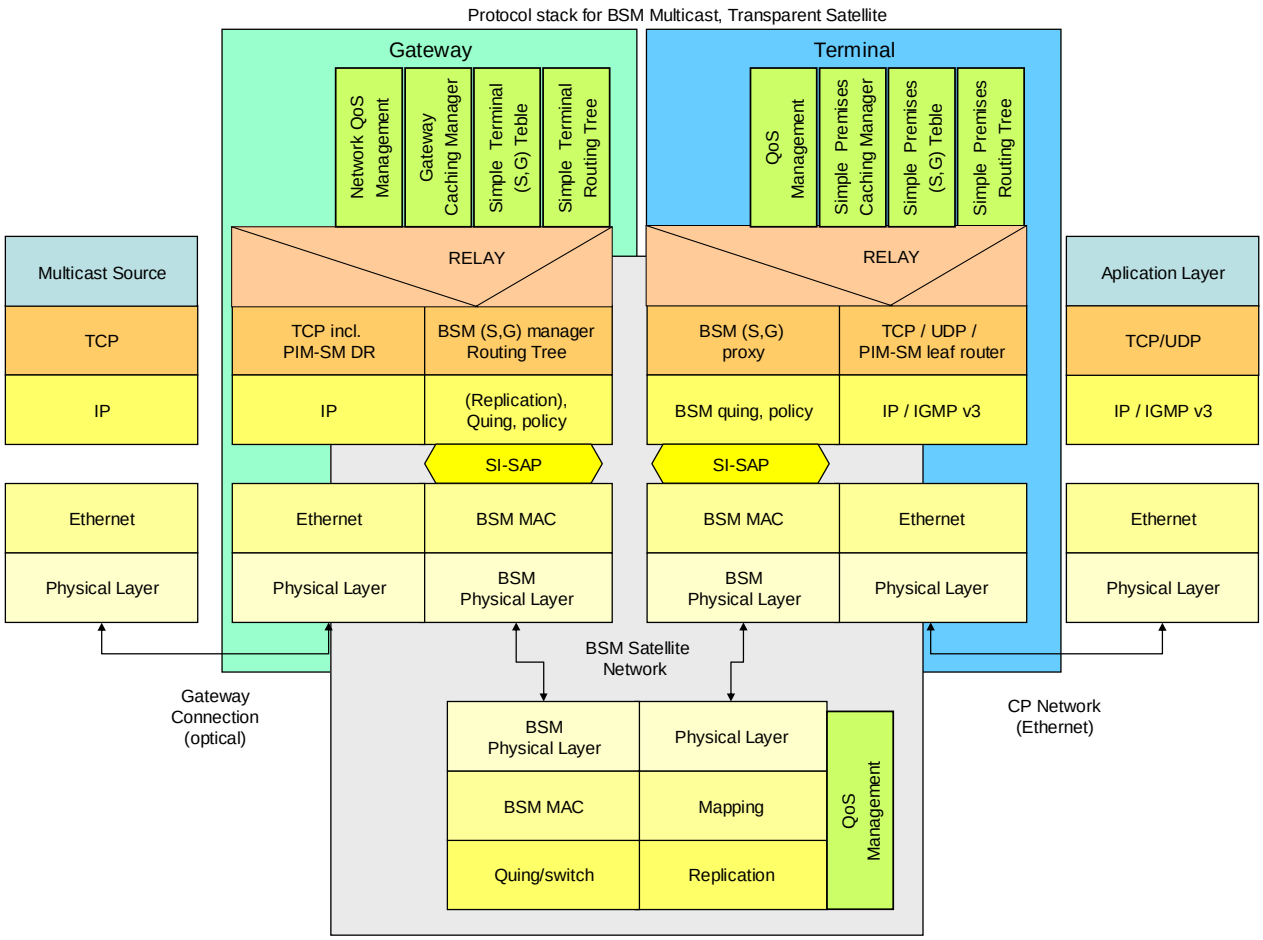

Figure 45: Draft sketch of a BSM Multicast protocol stack concept for OBP satellite

8.5.4 Performance measurement

Multicast performance measurement is usually straight-forward, as illustrated in IETF RFC 1889 [11] (Real-time Transport Protocol). A set of measurement hosts send small probe packets to a particular multicast session. It will also receive packets from the session in order to determine session transfer (network) performance.

Measurement metrics for a multicast session:

- Loss -percentage of packet loss from one client to another.
- Delay - one-way delay (in ms) from one client to another.
- Jitter - variation (in ms) of the one-way delay.
- Order - percentage of packets which arrived out-of-order.
- Duplicate - percentage of duplicate packets.

For a BSM network there must be individual measurements for source within and outside the satellite network. Measurements may also depend upon the number of gateways and satellite used.

8.5.5 Use case concepts

Technically, BSM-MP can work like this when it is at the edge of the network, and the satellite is a bent pipe type: A host connected to a terminal wishes to receive a source and subscribe to a multicast group. It informs the nearest router, residing in the terminal, which relays the message to its gateway.

When a host request a specific source, then the gateway that serves and manages the terminal will initially receive the source before sending it to the terminal(s). This is true both for external sources and internal sources from other terminals unless the satellite is OBP. At this point it can consider the requirements for the source and terminals (QoS setting, reliable multicast, capabilities of the terminal, subscription perhaps, and so on).

For non-interactive and non time critical services a gateway multicast manager could receive the multicast content, buffer it and adapt formatting to the BSM system before forwarding it over the satellite link to its destination host via any suitable protocol. It could also buffer it in case something needs to be resent. For real-time data it will be essential that received data is forwarded immediately to the end host, so that the interception of the data does not introduce any significant delay. There may be some buffering delay, but there is no logical fixed and unavoidable delay associated with the multicast manager approach (that is identified as of now).

The BSM-MP would in principle subscribe to all the relevant multicast groups on behalf of the hosts connected to the satellite terminals and guarantee that the content was delivered from the gateway multicast manager to the satellite terminal where it again would be forwarded to the destination hosts (or cached if relevant). The initiative to subscribe to a multicast group would come from the end user, but the multicast manager would receive the content (once) and then keep track of which terminals on its network that should receive copies of the content.

Possible disadvantages of the proposed concept is that it could introduce delay, it could intrude privacy, it would in principle break the protocol as satellite subscribers would not be the actual hosts seen from the outside. These things need to be handled when the specific standards are developed.

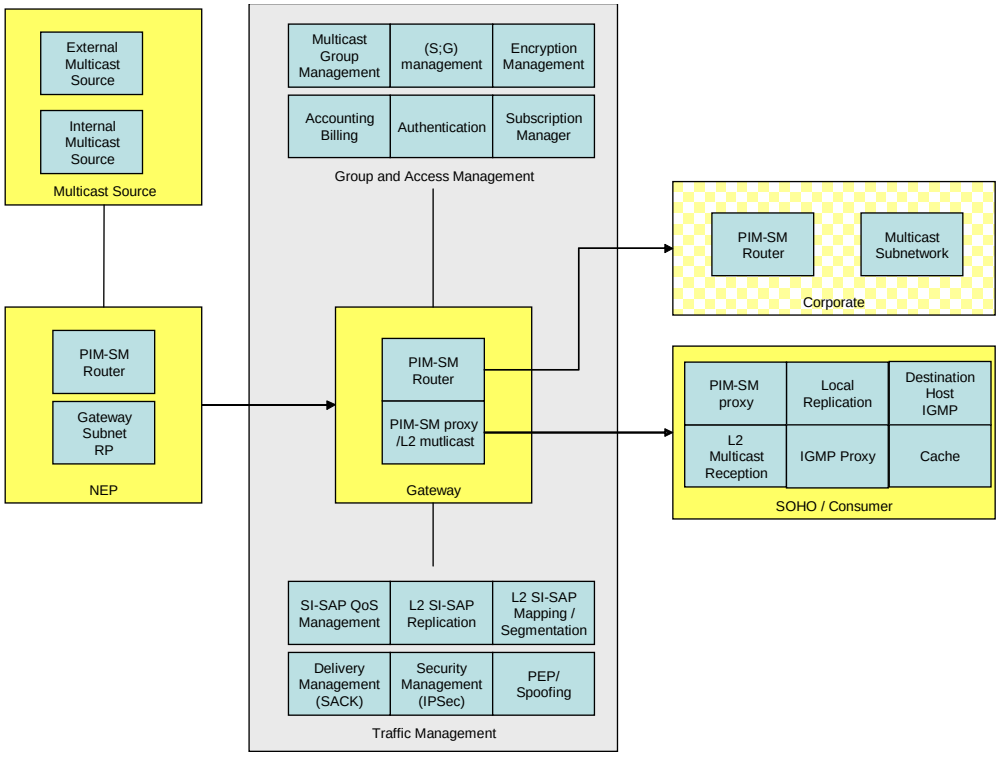

Figure 46: BSM Multicast functions

8.6 BSM multicast operation

Broadband satellite multimedia systems will all differ with each other with respect to satellite packet regeneration and supported network topologies, star, multi-satellite, mesh, etc. But the most interesting difference between BSM satellite systems affecting multicast may be the number of addressable downlink spot beams. A system with a single global coverage beam will differ considerably from a system which supports hundreds of spot beams, tens of regional beams, a global beam, a few simultaneous dynamically configured arbitrarily shaped beams and inter-satellite links as well. Multiple spot beams may be the best way to improve overall traffic handling capacity. However, the BSM system design should not inefficiently handle multicast traffic.

Therefore, given an arbitrary community of interest, one can see the advantage for a single source up-linking a minimum number (e.g. one) of packet streams and having the satellite replicate these streams into a configurable subset of these addressable beams (many).

In other words, the problem becomes intercepting the IGMP and PIM-SM IP signalling traffic at the BSM satellite terminals (acting as routers) and engaging in a new satellite dependent protocol to set up the distribution (replication) tree at L-2.

This clause discusses a common protocol, which should be able to have satellite independent elements. The goal would be a common ST state machine and SI-SAP messages. Information elements may be different. For example, one system may have a larger L-2 address space than another, and so on but in all cases there would still be a need for address resolution between the Class D multicast IP address and the L-2 multicast group ID. As far as possible, the protocol should "hide" the details of the satellite system from the IP layer. For example, on more complex systems, the satellite system may adopt different L2 services depending on the multicast service requirements (e.g. it may change the L2 functions from source distribution (multiple unicast), to satellite distribution (multicast) to single beam broadcast to optimize efficiency.

There may be other impacts on the IP protocols for multicast, such as delay and limited bandwidth. If the satellite delay may cause protocols to time out or if the signalling protocols are too verbose, then intercepting them and using the SI-SAP protocol (as possible subset of the larger protocol) may be beneficial to even the simplest of satellite systems. The effect of the satellite link on PIM-SM and IGMP needs to be studied, in fact.

8.6.1 General operation

BSM multicast operation can occur on demand or be scheduled by the source. For on-demand operation, multicast setup is triggered by data flow or customer premises protocol. For scheduled operation, multicast setup is based on time triggers configured by an NSP and NCC interaction. Multicast operation can also have static or dynamic group participation. For dynamic participation in multicast sessions, the BSM system proposes a receiver-based multicast since it is a more scalable form of multicast.

There are several aspects to supporting multicast over the BSM system:

1) Configuration of multicast.

2) Multicast setup over the BSM broadband system - scheduled or on-demand.

3) Multicast group management (i.e. IGMP (Internet Group Management Protocol) processing - supporting dynamic join and leave operations - recently approved IGMPv3).

4) Multicast routing protocols (e.g. Protocol Independent Multicast - Sparse Mode).

5) Multicast teardown.

A multicast group is a common interest group of parties that can participate in multicast sessions. A multicast session is a specific instance of multicast communication by a multicast group defined by its parameters (e.g. rate), participants, and time of existence. The source or origination ST is the root node of the multicast session and is an ST that can transmit data to a particular multicast session. A destination ST is a leaf node of the multicast session and is an ST that receives data from a particular multicast session. This protocol should not preclude the possibility of multiple root nodes in a multicast session.

An NSP can request that a multicast session among group members be set up beforehand. This is called a scheduled multicast session. A scheduled session may have participants that are static (preconfigured), or may allow potential participants to join and leave dynamically after the session has been set up.

A multicast session can also be requested on-demand as a result of data flow or signalling. This is called an on-demand session. An on-demand session may have participants that are preconfigured or may allow potential participants to join dynamically after the session has been set up.

End hosts use IGMP as the standard protocol for letting routers that are on their local network know that they are interested in receiving data for certain multicast sessions. If the ST is the Designated Router for the end host, the ST will receive the IGMP join request. Based on receiving the IGMP request, the receiving ST sends its address information to the NCC. The NCC can admit the join request and, if necessary, request an update to the distribution tree to ensure that multicast data arriving at the payload is replicated and sent to this ST. The ST then forwards the data to the port of the requesting host. Note that subsequent IGMP requests from other ports do not require a new Dynamic Join request with the NCC; the ST should locally replicate data to the port that joined, if allowed by configuration.

Multicast session announcements and session detail information (e.g. posted to a website) are standard means for distributing information relating to multicast sessions to actual or potential participants. Multicast session announcements may occur outside the scope of BSM.

Multicast routing protocols are needed in order to enable multicast forwarding of packets addressed to members in the multicast group. Since IP multicasting allows group members to join or leave a host group at any time, the topology of a group's multicast delivery tree can change and the routing protocols keep track of those changes. Data is forwarded only to those satellite terminals (ST) that have multicast members connected to them.

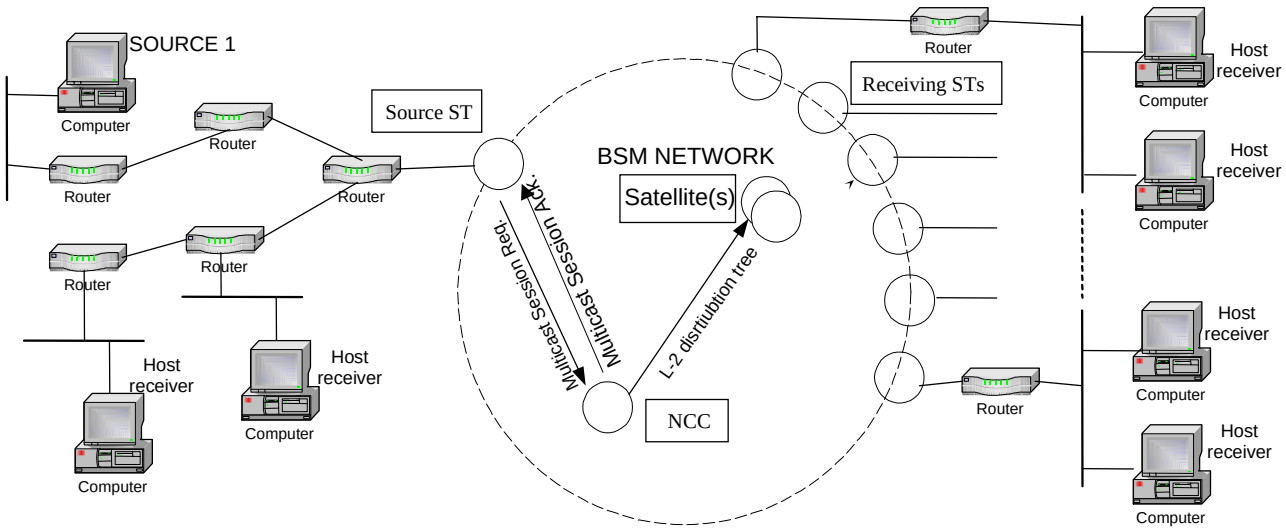

Figure 47: BSM multicast concept with drawing details

8.6.2 Preconditions and Assumptions

The following preconditions are assumed for the multicast scenarios:

- Security has provided the appropriate authorization information to the ST from the NCC for the multicast. Note that conditional access and key distribution are beyond the scope of the present document.
- The coordination of senders and receivers in a multicast is managed by applications outside of the BSM system.
- We show message flow directly between the NCC and the satellite independent layer. In fact, one primitive may flow between the satellite dependent layer and the satellite independent layer. And, in turn, this primitive may be mapped to a system dependent message between the ST and the NCC.

8.6.3 Scenario description

8.6.3.1 Setup of scheduled multicast sessions

The Network Service Provider (NSP) is the owner of multicast sessions that are scheduled for satellite terminals (ST) participation. The NSP provides the NCC with the session details and the NCC uses those details to determine whether or not to set up the multicast session. This scenario is shown below. After the NSP has set up the multicast, a preconfigured session or dynamic join can take place, depending on the type of session.

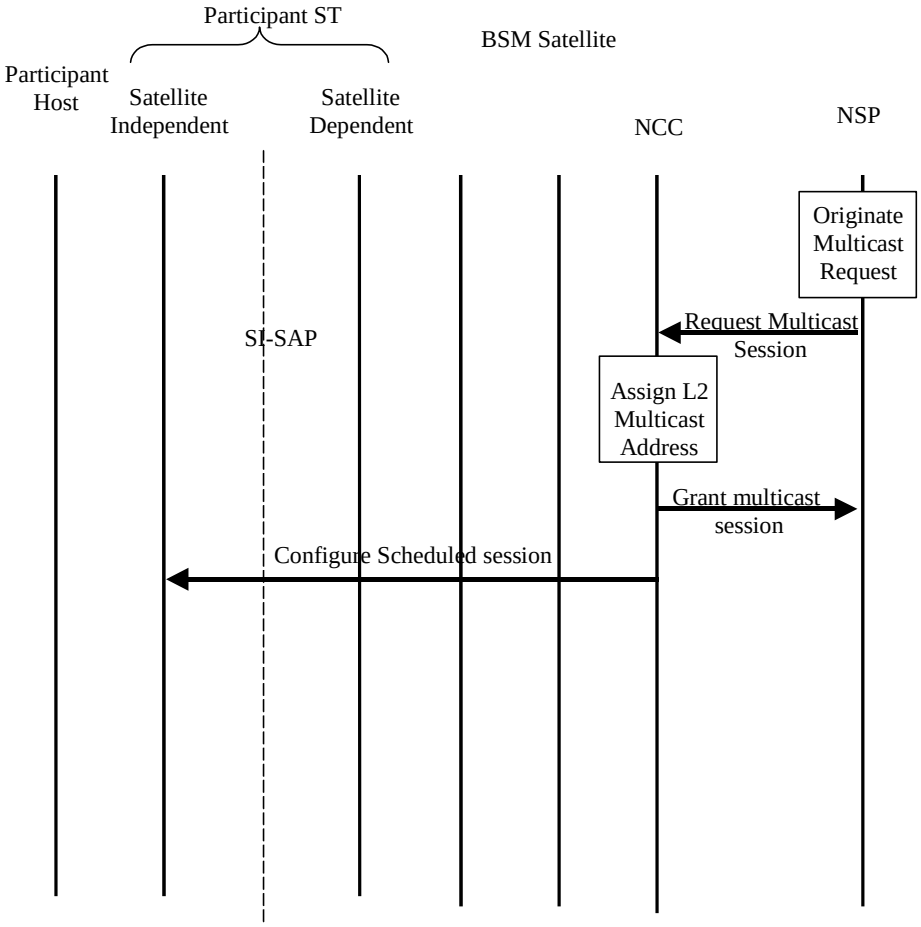

Figure 48: Setup of scheduled multicast sessions scenario

1) NSP management checks multicast session requests against policy and current state.

2) NSP sends a request to establish a multicast session in the future. The request contains the information on the rate, the Class D multicast IP address, the source ST, and times of the multicast session. If this multicast session is also going to be preconfigured, this information can include the list of STs that will participate in the session as well. For dynamic membership sessions, the potential ST participants may be included.

3) System management at the NCC assigns a layer 2 Multicast Address for the session. The layer 2 Multicast Address maps to a Class D multicast IP address.

4) System management at the NCC sends the NSP a confirmation that the multicast session has been scheduled.

5) System management at the NCC sends to all the known participating STs the configuration information for the scheduled session including the start time, duration, Class D multicast IP address, layer 2 Multicast Address, and rate. Each source ST receives the additional information required to establish the multicast session at the time of the scheduled multicast. If necessary, new classifier rules are also provided to the source ST to allow it to map incoming user data into the scheduled multicast session. In general, there may be satellite dependent as well as satellite independent configuration.

8.6.3.2 Setup of on-demand multicast sessions

This scenario shows an example of how a multicast session can be set up on demand by an end user. The on-demand multicast session can be stimulated by IP data with a multicast class IP address or by a signalling request to set up multicast. The origination ST must be configured with the appropriate classification rules before it can launch a multicast session setup request to the NCC.

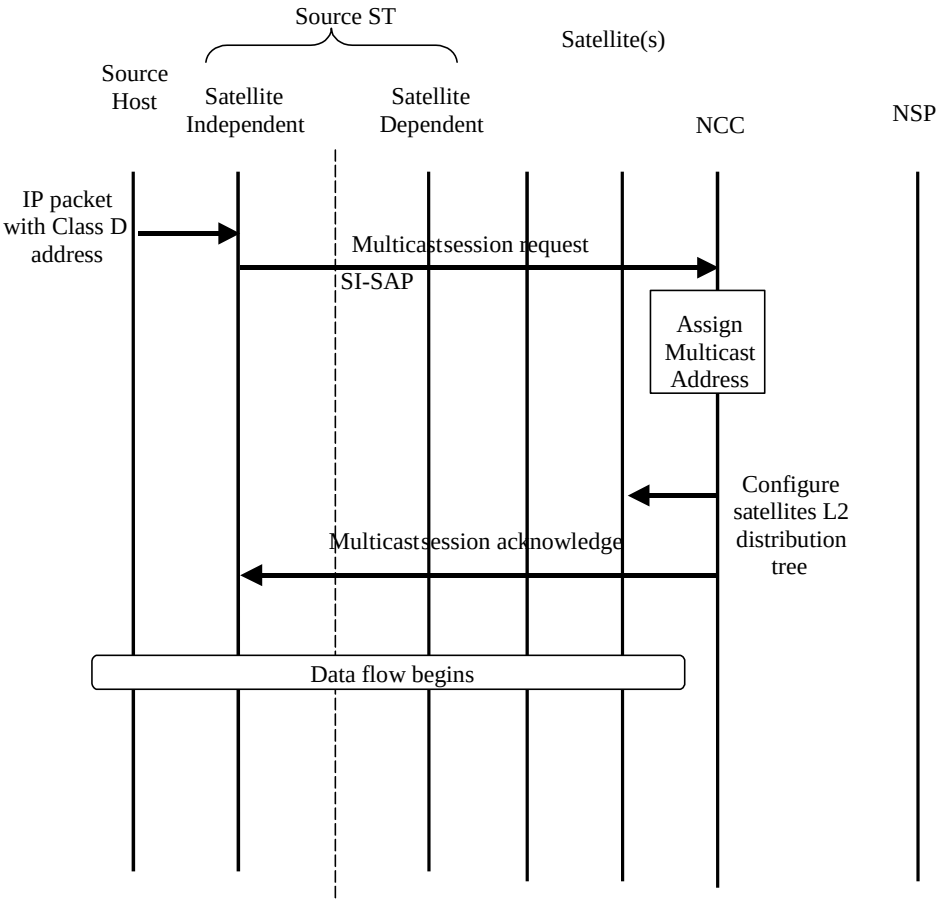

Figure 49: Setup of on-demand multicast session scenario

1) A packet arrives from the origination host that contains a stimulus to launch a multicast session. In this scenario, an example of a stimulus is an IP Packet that contains a Class D multicast IP address. This stimulus may be any user data including H.323, SIP, or other multicast protocols.

2) The ST sends a multicast session setup request to the NCC including information such as desired rate and IP multicast address.

3) The NCC determines whether this source ST has the ability to initiate a multicast session with the requested parameters. The NCC checks the request against its current view of system resources and capacity.

4) If the multicast uses satellite packet replication, the NCC sends the required information to the BSM satellite(s). If no leaf users have yet to join the multicast session, the step of assigning the Multicast Group ID and updating the payload may be delayed.

5) Multicast session management procedures are completed to establish the multicast session. For cases where there are no ST receivers who have yet joined a dynamic on-demand session, the NCC sends the source ST a Progress message that starts a longer timer cycle. The ST either times out if no receivers join or it receives a session established message that contains the destination information required to address multicast packets for transmission for the receiver(s) that joined.

6) User Data Transport for the multicast session can begin.

8.6.3.3 Dynamic join

Figure 50 shows one of several cases of multicast setup when the receivers are joining dynamically. A multicast session is already active when the destination host requests to join the multicast session. When an end host requests a multicast via IGMP, the ST at that end receives the IGMP or a multicast routing protocol join message and makes a request to the NCC to join the multicast. The NCC, based on the request, generates an update to the distribution tree, if necessary. The NCC then grants permission for the receiving ST to join the multicast. The receiving ST may start to receive data immediately.

An ST that is not attached directly to the LAN of the requesting host will not receive an IGMP message, but will receive a multicast routing protocol message (e.g. PIM-SM join) if the routers along the path support multicast. The ST terminates the PIM-SM message or IGMP message and, if configured to do so, creates a message to the NCC. The NCC determines if this is the first join message of the multicast or the last prune message of the multicast within the BSM system. If so, the NCC sends a message to the source ST that prompts the source ST to generate a PIM-SM join or prune message towards the Rendezvous Point (RP) using the unicast routing information. An ST may support the functionality of acting as a Candidate RP or a BootStrap Router (BSR).

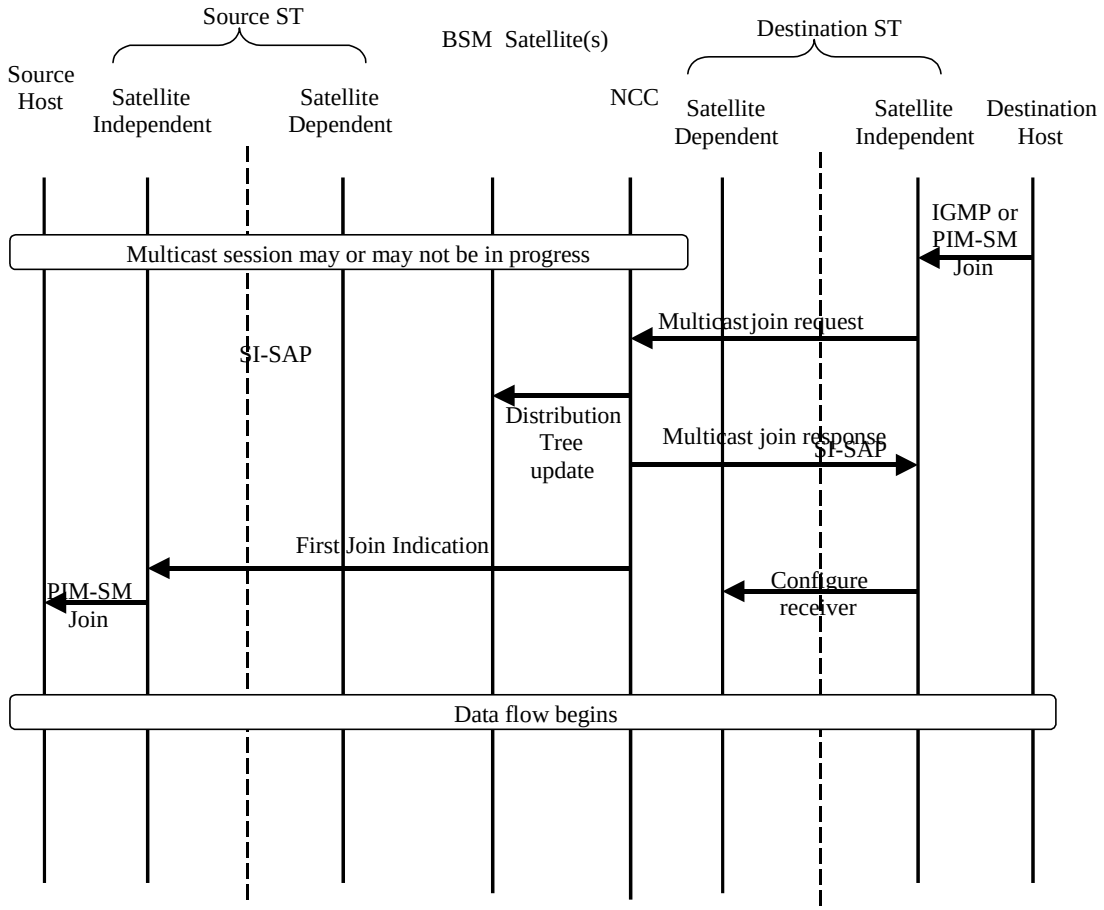

Figure 50: Dynamic join

1) A multicast session may exist with an assigned layer 2 Multicast Address. The destination ST has previously been configured with the Class D Multicast IP address of the multicast session along with the list of eligible ports. The destination host has previously received details such as the Class D Multicast IP address of the multicast session from an application level mechanism (e.g. web page). The session may or may not be active when the host attempts to join.

2) The destination host sends an IGMP request towards the ST to request participation in the multicast session. The ST receives either the IGMP message or, if a multicast-capable router exists between the host and the ST, the ST receives a PIM-SM join message.

3) The request is checked locally against the ST configuration (e.g. eligible port list) to determine whether the ST has permission to request participation in the session and the ST sends a multicast join request to the NCC.

4) The NCC determines whether this destination ST has the ability to participate in this particular multicast session (e.g. checks user group restrictions). The NCC updates the distribution tree if required.

5) The NCC responds to the destination ST.

6) Upon receiving the grant to join the session, the destination ST configures its receiver to accept packets destined for the layer 2 Multicast Address.

7) If this is the first ST in the BSM network to join the multicast, the NCC also sends a message to the source ST port in order to prompt the ST to generate a PIM-SM join message to build the multicast tree.

8) The source ST sends the appropriate PIM-SM join message upstream towards the Rendezvous Point (RP). The RP may be directly connected to the ST or the message may need to traverse through one or more multicast-capable routers on the terrestrial network.

9) For an existing session, the origination ST data is also forwarded to the new participant as well as all existing participants.

8.6.3.4 Dynamic leave

For sessions that allow dynamic membership, any participant may leave the session while it is active. Upon leaving the session, the current multicast configuration dictates whether or not this affects the distribution tree. If other STs in the multicast session are still active in the satellite beam, for example, multicast packets are still down-linked to that beam. In other words, the satellite replication table need not be updated. However, the satellite dependent receiver of the ST is reconfigured so that it no longer receives packets for the session.

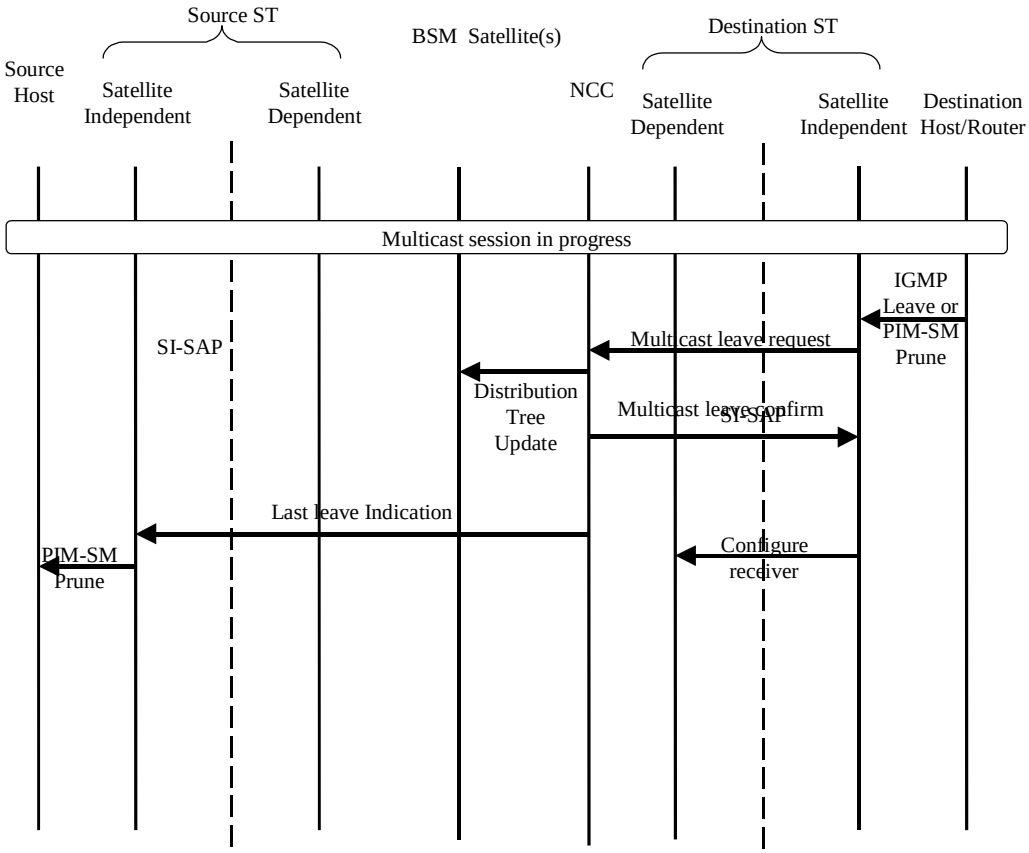

Figure 51: Dynamic leave of multicast session

1) An active multicast session exists. The destination is receiving data from the multicast session.

2) The ST receives either the IGMP Leave request or a PIM-SM prune request (if one or more multicast-capable routers are between the host and the ST).

3) Assuming that this is the only host that the ST serves that is participating in the multicast session, the ST launches a Multicast Leave request.

4) The NCC determines whether the removal of the ST from the multicast changes the L-2 distribution tree.

5) The NCC confirms the leave to the ST.

6) The ST reconfigures the satellite dependent receiver to no longer receive packets addressed to the Multicast Group ID for this multicast session.

7) If this is the last ST in the BSM system to leave the multicast, the NCC also sends a message to the source ST port in order to prompt the ST to generate a PIM-SM prune message upstream out its terrestrial port to teardown the multicast tree.

8) The source ST sends the appropriate PIM-SM prune message upstream towards the Rendezvous Point (RP). The RP may be directly connected to the ST or the message may need to traverse through one or more multicast-capable routers on the terrestrial network.

8.6.3.5 Multicast teardown

A multicast session can be terminated either by reaching the end of a scheduled duration, by signalling from the session owner that the multicast session has ended, or, for sessions that have no duration, when no data has been sent for a configurable amount of time. The NCC signals the release of the multicast to all the participating STs.

8.6.3.6 Configuration parameters

For on-demand multicast sessions, the following parameters should be configured in all potential source STs:

- Rate.
- Class D Multicast ID.
- Multicast Group ID.
- Classification Rules.

For scheduled multicast sessions, a time profile (start time, duration, and schedule) should also be configured in source ST:

For membership in a multicast session, the following parameters should be configured in destination STs:

- Class D Multicast Address.
- Multicast Group ID.

9 Recommendations

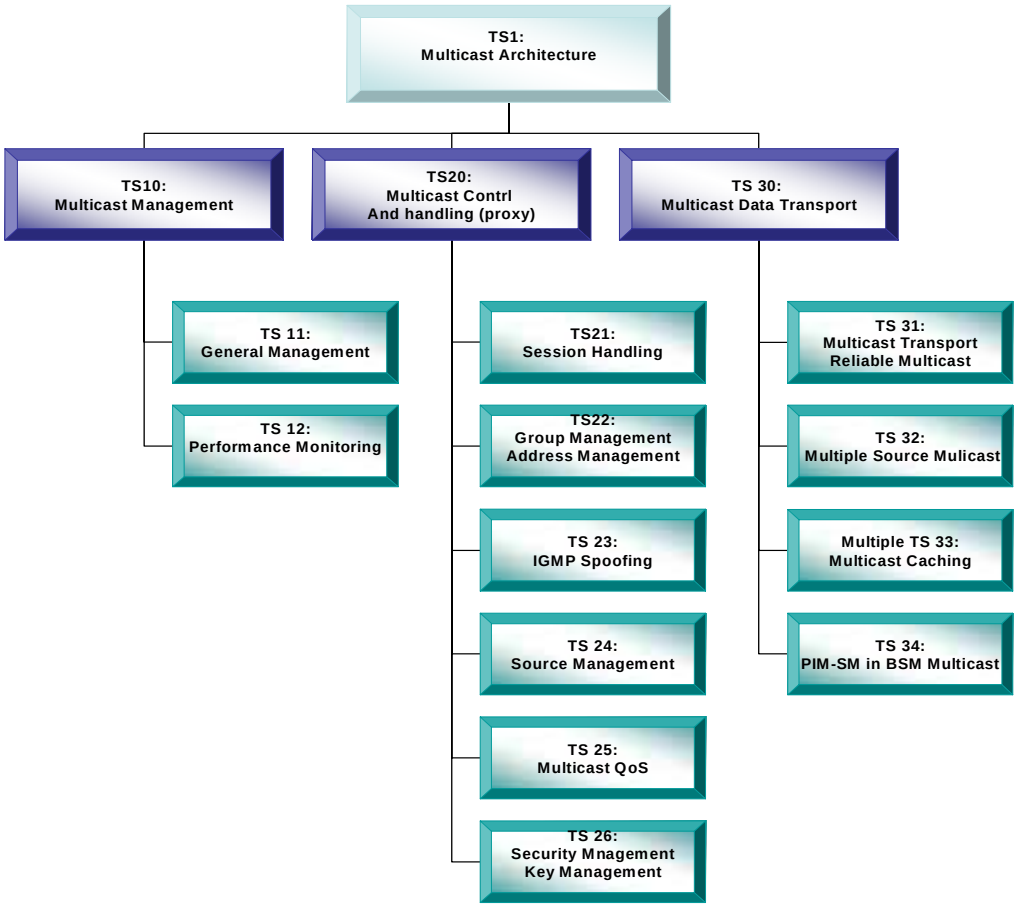

Figure 52: BSM multicast specifications sketch

9.1 Specification topics and draft scope

The following set of **specifications topics** are proposed, with a draft scope given below each one. It may not be essential to prepare the specifications as individual document, as they can also form separate chapters/volumes of a multi-part BSM multicast specification.

The topics' priority is to be considered outside the present document.

- **TS 1: BSM multicast architectures for OBP and bent pipe**
  - Functional specification.
  - There will be a few (2-3) different multicast network architectures for BSM networks depending on the number and type of satellites, spot-beams and gateway network structure.
  - The TS will not list all possible BSM multicast architectures. More reference architectures can be added later as required.
  - The TS will define architectures, which will be used in the further standardization work.
  - Initially the BSM families should be supported, as well as other ETSI standards like EN 301 790 [34].
  - The TS will also analyze the concept of distributed sub-networks under BSM multicast networks.
  - Multicast architectures will likely be based on figures from the present document.
- **TS 10: Multicast management and monitoring**
  - Contains a summary and overall description topic and of the documents on the tree below.
- **TS 11: General multicast management in BSM systems**
  - The scope of this TS will be to define how the overall multicast process is managed, including issues that are defined in FCAPS (fault, configuration, authentication, performance, security) management.
  - It may be essential that BSM multicast management be managed similarly as other multicast technologies, as several major service providers may not wish to train a special group of people to handle the management of specialized satellite systems.
  - Definition of the ability to manage the multicast process and place operational procedures to deny/allow sources to use specific uplink terminals, and possibly to control their transmit rate.
- **TS 12: Performance monitoring on BSM multicast**
  - This TS will define which parameters should be monitored and how, in order to define and decide the performance of BSM multicast systems. This could possible form part of a TS on QoS, but as QoS in multicast is not yet in common use, and a performance evaluation method for BSM multicast will be from day one, it is probably best to develop this TS early. If beneficial, it could later form part of another document.
- **TS 20: Multicast handling and control (proxy)**
  - This family of documents describe the basic functionality needed to administer and set up data transport and sessions.
  - The top document gives an overall description.
- **TS 21: BSM multicast session handling and management**
  - The session management service needs to advertise scheduled sessions and provide a query mechanism for retrieving information about session schedules.
  - Connection management.

- **TS 22: BSM multicast group and address management**

  - The TS will define suitable use of IGMP, SSM, PIM-SM and other RFCs in BSM networks. Focus will be on identifying performance issues, for instance related to IGMP version 2 versus 3.

    - *With SSM the Network Operator has a trivial address allocation, with 16 million addresses per host, there is no network-layer source discovery (PIM moved to the application layer) and the content Provider can be given exclusive access to specific multicast groups. Permanent multicast groups are easy to advertise. In a multiple gateway scenario SSM can be used with SSM Proxies, where the sender (co-located in the NCC or in a gateway for instance) unicast to nearest proxy (in a gateway). If required a proxy then relays data to other proxies, but finally all proxies multicast to attached receivers. Receiver allocation to an SSM proxy can be static or dynamic. ASM is not recommended.*

  - PIM can and will be used when the satellite network is not an edge network.

  - The TS will define procedures on how to interact between BSM network and the above IETF protocols.

  - Define how BSM layer 2 address space relate to layer 3 IP address space.

  - The TS should specifically consider a concept that is reflected also in a recent Internet draft (Bill Fenner et al. "draft-ietf-magma-igmp-proxy-01.txt", July 2002) [38].

    - *In certain topologies, it is not necessary to run a multicast routing protocol. It is sufficient to learn and proxy group membership information and simply forward based upon that information. This draft describes a mechanism for forwarding based solely upon IGMP membership information. This document applies spanning tree multicast routing to an IGMP-only environment. The topology is limited to a tree, since we specify no protocol to build a spanning tree over a more complex topology. The root of the tree is assumed to be connected to a wider multicast infrastructure.*

- **TS 23: IGMP spoofing**

  - This TS will define how to minimize the administrative overhead with running IGMP over satellites in BSM networks.

- **TS 24: BSM multicast source management**

  - This item will define how a multicast source may "connect" to the BSM network, including optional local sources.

  - It will consider knowledge of where the source enters the BSM network, where it comes from and where in the network it is to be delivered. This may depend on the source address, or contributing network.

  - Handling of sources for closed groups, like corporate multicast sources are to be defined.

  - Handling of multicast originated within the BSM network is going to be defined. This will clearly depend upon the satellite capability, and the network structure.

  - Time to live or scoped multicast will be defined.

  - The TS will also define how hosts subscribe to sources, and possibly also define restrictions on what categories of terminals (capabilities, subscription or ownership) may receive (or transmit) which sources.

  - The TS may define some basic security requirements, but the actual security (algorithms, encryption, etc.) is best handled in a separate TS.

- **TS 25: Multicast performance and QoS**

    - *QoS settings need to be able to interoperate in the outside world, but also be able to provide particular benefits for satellite system customers. Thus a BSM multicast QoS management system should be able to take advantage of the particular features of a satellite network. It will commonly be possible to adapt coding, and sometimes modulation, to the individual users. In the case of multicast, to the multicast data and to the type of data.*

  - Negotiation of QoS settings between source/gateway/sender and terminal/receivers should be defined.

- Also QoS monitoring schemes and the numerical definitions should be defined and decided.
- Procedures should be defined on how to handle cases when QoS settings cannot be fulfilled temporarily for instance due to rain fades at Ka-band.
- QoS in multicast adds a new dimension if both content and receivers can define required and desired settings.

- **TS 26, TS 26.1, TS 26.2: Multicast security**
  - Limiting Senders, Limiting Access, Limiting Receivers. This TS should focus on how data can be secured so that nobody who is not in a multicast group can and will receive it. This will be important for corporate multicasts on a shared satellite network, as opposed to a traditional corporate LAN, where there can be firewalls and other security means.
  - Verifying Content as discussed in IETF RFC 3170 can be handled.
  - The TS will also focus on how to ensure the multicast source does not get jammed by spam senders or hostile hackers.
  - The TS will defined a secure group member management scheme with a password check or other means of identification, authentications and authorization before groups can be joined (or perhaps even left).
  - Protecting Receiver Privacy.
  - The TS will defined how secure data can/shall be supported in BSM multicast networks.
- **TS 30: Multicast data transport**
  - This set of documents describe how the actual transport of multicast content is taking place.
  - The top-level document will present an overview of the concepts involved.
  - It will be updated as the documents below evolve.
- **TS 31: Multicast transport and reliable transport in BSM networks**
  - The TS will define segmentation and replication.
  - Reliable multicast transport.
  - The TS will describe BSM multicast interfaces to the Internet and probably define basic primitives.
  - The TS will also define efficient data delivery, considering both broadcast type of systems like DVB-S and other BSM family systems with multiple spot-beams, since multicast data should only require capacity in those beams that have subscribers to the content.
  - If the content requires reliable multicast as opposed to best effort, then the underlying satellite multicast system should define issues like the use coding, modulation, which can differ in every beam, and possible the use of ARQ protocols. QoS issues may come into play here (possibly in a future revision of the TS).
    - *When the BSM network is an edge network, which it is in many cases, then UDP/IP/BSM can be integrated and optimized together. Mapping the messages optimally onto carriers/timeslots is necessary to conserve bandwidth spectrum and the minimum amount of replication requires some control over the RRM functions. Terminals in a beam should be able to listen to the right carriers and timeslots. Further, terminals with more than one (but a limited number) host receiver for a multicast message should be able to replicate the messages locally. This is relevant for SOHO networks and residential networks. ITU is also working on a new L3 protocol, and it is worthwhile to consider collaborating on this. Developing a new transport protocol requires inputs from companies and research bodies.*
  - The TS will define BSM Multicast Addressing for the architectures under consideration.

- **TS 32: Multiple source multicast**
  - The scope is to define how multi-party multicast, like in multiparty video-conferencing, shall work on BSM networks.
- **TS 33: Multicast caching**
  - These TS will define the concepts of BSM edge- and data casting over multicast satellite systems. Caching of data can be either at the gateway or at the terminal or both.
  - The TS will define user control parameters and network operator parameters. For instance if the users technically owns the terminal and the storage device, what permissions shall be granted for 3rd party use of the hard drive for potential useful content.
  - Security of cached content should also be treated.
- **TS 34: PIM-SM in BSM multicast**
  - Assuming general IGMP over satellite, this document will describe how PIM-SM also can be implemented and integrated into the system in an efficient manner.
  - Assume only few subscribers require pure PIM-SM. Majority will run IGMP.

Things ETSI should not do include:

- Defining "competing" multicast protocols with IETF standards.
- Defining standards locked to specific air interfaces, possibly with the DVB exception.
- Defining standards that are out of line with technology trends and business models.

9.2 Brief discussion on concept

The above recommendations indicate that ETSI should develop a (set of) BSM multicast standard(s) that is based on using a multicast proxy manager concept.

There are benefits in having a standard way of how a BSM satellite system interacts with the outside world, both from the terminal and from the gateway side. However, internally the system should allow for freedom in technical design, that can take individual BSM satellite system characteristics into account. The multicast manager should therefore take advantage of the SI-SAP work performed in the ETSI BSM workgroup. The layer 1 could be any BSM family, and the Multicast Manager should interface over the SI-SAP. The multicast manager will in general interface at Level 3 or above, but commonly operate internally at layer 2 or 3 or at a combined L2/L3 level.

The multicast proxy manager would not be fully transparent for IGMP multicast, in the sense that it would aim to conserve bandwidth and adapt messages and overhead to the available capacity in the BSM system, yet be compatible with the IGMP interface on the terminal side. The BSM system should therefore be able to spoof timers and management overhead. Also added security mechanisms can be provided that prevent hacking/jamming or other security threats.

Some large sub-networks like corporate networks with PIM-SM routers may need to interface to other PIM-SM routers. Transparent functionality (for PIM-SM) may therefore need to be present in the system, and to made available for large sub-networks, i.e. when the satellite access functions acts as a local gateway.

The specifications will also allow for adding local multicast groups at the multicast manager level. Local multicast groups can have a QoS control, even when this is not supported over the external network.

The BSM multicast specifications should treat the satellite network as an independent subnetwork, and often an edge network assumption is also valid, where it interfaces to the external (edge of the) core network via a given satellite connection point. A logical entity (like the MEP) will be able to provide information on the gateway network structure, or on which gateway (or BSM system entry point) that is to be used to reach a given satellite terminal (that in turn can connect several hosts). Thus, terrestrial optimal routing algorithms (not considered here) may be implemented, minimizing the terrestrial travel path or some other cost function.

If more than one host subscribes to multicast content then the satellite terminals should (will have to) include a small multicast router (or another entity) that replicates the messages locally. The equipment does not have to be a full scale router to do this type of multicast forwarding, as the anticipated number of connected hosts will be small, in the order of a fem less than ten, hosts (for a domestic home network).

A rough and basic illustration of some of the functions is offered below, but this diagram is only meant as a draft sketch, and is subject to change as work progresses.

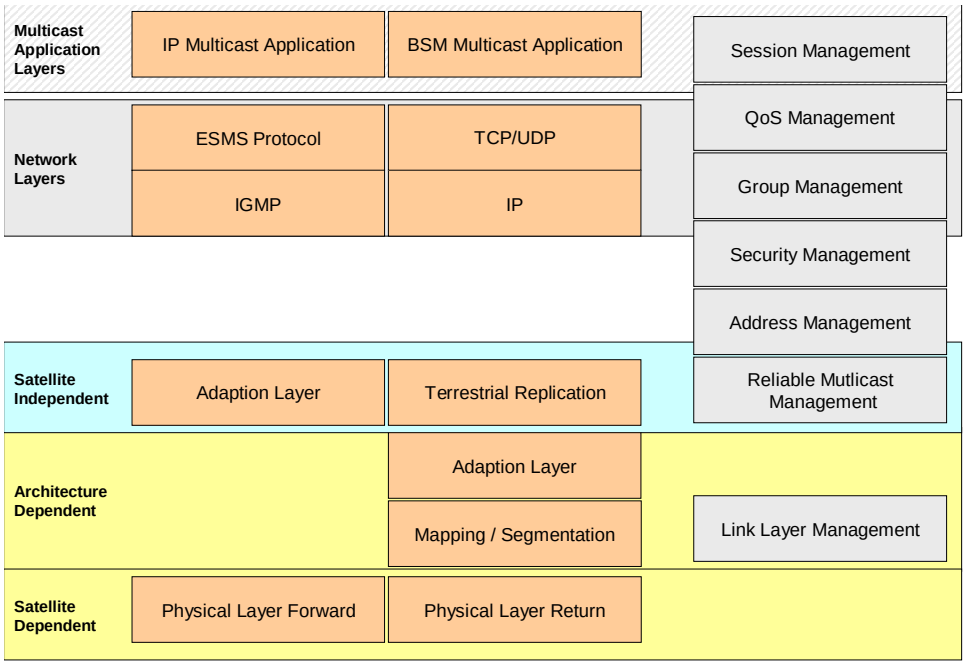

Figure 53: BSM multicast protocol layers and management functions

Annex A:
Security systems in DVB-S and DVB-RCS

Security in general is intended to protect the user identity including its exact location, the signalling traffic to and from the user, data traffic to and from the user and the operator/user against use of the network without appropriate authority and subscription. In DVB, two levels of security can be applied:

- DVB common scrambling;
- Individual user scrambling in the forward and return links.

In addition, security can be applied in the application, transport and network layers. Application and transport layer security are not discussed here.

Although the user/service provider could use its own security systems above the data link layer, it may be desirable to provide a security system at the data link layer so that the satellite link is secure without recourse to additional measures. Link level security is particularly desirable by satellite access network operators in order to secure satellite links and provide their clients (such as ISPs) with data confidentiality.

For DVB, the satellite interactive network forward link is based on the DVB/MPEG-TS Standard. The security concept is shown below.

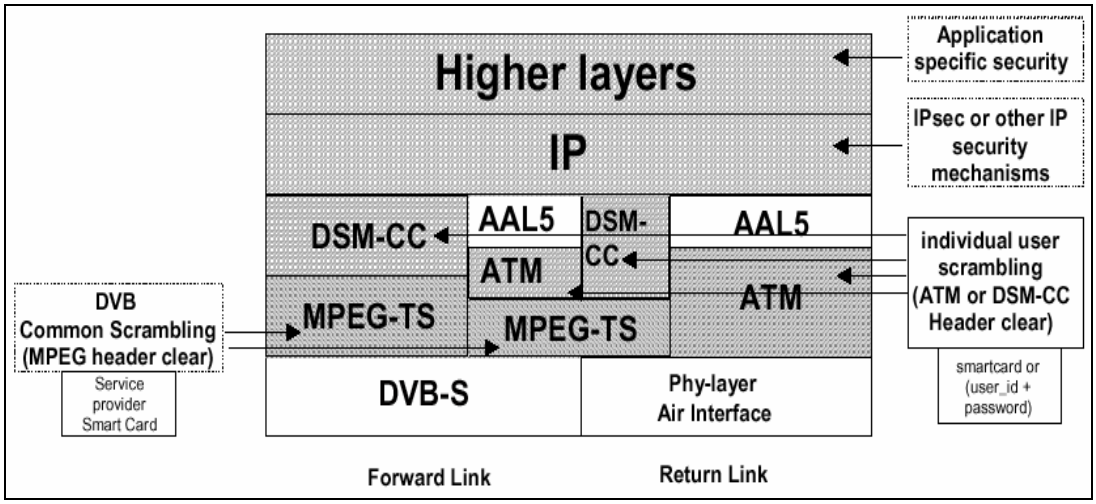

Figure A.1: IP stack in DVB-RCS

A.1 Conditional access in DVB-S

Conditional Access (CA) is a service that allows broadcasters to restrict certain programming products to certain viewers. The CA does this by encrypting the broadcaster's programs. Consequently, the programs must be decrypted at the receiving end before they can be decoded for viewing.

CA offers capabilities such as Pay TV (PTV), interactive features such as Video-On-Demand (VOD) and games, the ability to restrict access to certain material (such as movies) and the ability to direct messages to specific set-top boxes (perhaps based on geographic region).

DVB Conditional Access (CA) originated as a broadcast security mechanism that allows a source to determine which individual receivers are able to receive particular broadcast programs. CA requires two principal functions: (a) the ability to encode (or "scramble") a transmission and decode it (or "descramble") at the receiver, and (b) the ability to specify which receivers are capable of descrambling the transmission.

The transmission from a source to all receivers comprises a set of scrambled MPEG components (video, audio, data); Entitlement Control Messages (ECMs, session keys); and Entitlement Management Messages (EMMs, service keys). The ECMs identify the CA services, and for each CA service carry the Control Word (CW), in an encrypted form, and any other parameters required accessing the service. The Entitlement Management Messages (EMM) are a set of messages that identify the entitlements (permissions) of any individual user.

In addition, a Subscriber Management System (SMS) maintains and stores commercial aspects of customer relationship (registration, granting of entitlements, invoicing, and accounting), and the Subscriber Authorization System (SAS) encrypts codewords and delivers them to the descrambler.

At the receiving end, it is the job of the Set-Top Box (STB) to descramble the CA encryption and decode the MPEG-2 streams for viewing. Each packet has associated with it (in its header) a Program IDentifier (PID). The Conditional Access Table (CAT) has a well-known PID value = 1. This table can be used to identify the PID values of the transport packets containing the EMMs. The demux processor also constructs the Program Map Table (PMT) from non-encrypted packets; this gives the PID values of all the transport streams associated with a particular program. Private data associated with the program can also be included in this table - for example, the PID value of the packets that contain ECMs. All these tables (signalling messages) are transmitted in the clear, which is an inherent security weakness in DVB-S systems.

A.2 DVB-RCS security

The DVB-RCS standard provides much more advanced security procedures (in comparison to DVB-S CA) for satellite terminal authentication and key exchanges with the Network Control Center (NCC).

DVB-RCS security can be divided into two phases: Phase 1 is the authentication during the logon procedure. During this phase a security session key is agreed between the Satellite Interactive Terminal (SIT) and the NCC. In phase 2, the session key is used for the encryption of all subsequent messages between SIT and NCC. The authentication is based on a long-term secret shared between NCC and SIT, called a cookie. The cookie is 160 bits long and stored in non-volatile storage (such as smart cards). The NCC maintains a database of the cookie values of the SITs on its network. Cookie values can be updated occasionally as dictated by security policy, but they are less vulnerable than session keys. Anti-cloning measures can also be implemented using message sequence numbering. The DVB-RCS standard allows a Quick, Explicit and Main key exchanges.

In summary, the main messages during log-on are as follows:

- Logon: The SIT indicates its intention to connect to the satellite network.
- Security sign-on: The NCC indicates which cryptographic algorithms it supports, as the initial stage of a security negotiation.
- Security sign-on response: The SIT responds by specifying the specific algorithms and parameters it will use, chosen from the list presented by the NCC.
- Main key exchange: This message and the following enable the NCC and SIT to use a public key algorithm to agree a shared secret.
- Main key exchange response: The second message enables the parameter values of the public key algorithm to be calculated.

A further consideration is security of the space segment. In satellite systems with DVB on-board switching, message integrity between the NCC and the OBP is important in order to make sure that configuration messages originate from the NCC. The major constraint in the OBP is its limited memory and computational power, since the computational cost of message integrity can be high. This depends on the type of algorithms used. For example, message integrity can be provided using public-key digital signatures, which are computationally heavy, or using MAC (Message Authentication Code) with secret keys, which is lighter. The use of secret keys implies the need for a key agreement, where keys can be stored in the OBP at installation time or agreed using the DVB-RCS key exchange mechanisms.

A.3 DVB-S and IP multicast security

DVB-S conditional access is used today for digital broadcasting over satellite and can be used to secure multicast communications over satellites at the MPEG-TS level. In DVB-S, IP packets are encapsulated in an Ethernet style header called Multi Protocol Encapsulation (MPE), where the IP address can be associated with the MPEG-TS PID. IP multicast can also be encapsulated with MPE. Descrambling in DVB-S is program based, where a whole program will be scrambled with the same CW. The program may contain video, audio and data, each with a specific PID. The main drawback is that DVB-S scrambling system favours a centralized ECM and EMM and its use for securing dynamically changing IP multicast groups is limited.

On the other hand, the DVB-RCS standard provides more advanced security procedures for satellite terminal authentication and key exchanges with the satellite network operator. However it does not provide security procedures for terminal-to-terminal communications. The DVB-RCS standard allows the use of ATM cell transmission over satellites. Hence for satellite ATM networks, terminal-to-terminal communications and multicasting can be secured using the ATM security system.

A.4 Satellite ATM security systems

A.4.1 Technical challenges in GEO satellites

ATM security, as defined by the ATM Forum Security Working Group, is modelled after the ATM protocol reference model, which is divided into three planes: user, control, and management. The ATM Forum security specification applies to Virtual Channel Connections (VCCs) and Virtual Path Connections (VPCs) for both point-to-point and point-to-multipoint connections. The ATM Forum defines the support of the following security services in the user plane:

- Entity authentication.
- Key exchange.
- Data confidentiality.
- Data integrity.
- Access control.

According to the ATM security specifications either the two-way or three-way Security Message Exchange (SME) protocols may be used to establish the above mentioned security services. These SMEs can either be signalling or in-band based. Security negotiation parameters can only be exchanged using the three-way SME. For unicast connections, either the three-way SME or two-way SME can be used to set up security associations. For the first "leaf" of a multicast connection, again, either the three-way or two-way SME can be used; for subsequent leaves, only the two-way SME can be used.

The ATM Forum security specifications state that for the data confidentiality service the ATM cell-level approach is used to encrypt the payload, and the header is left in the clear. The data integrity service is provided at the AAL level (rather than the ATM layer). Once a connection is established, keys for integrity and confidentiality services are negotiated using the three-way or two-way SME. However, when a key is used to provide confidentiality and integrity protection, the probability of successfully "cracking" the key increases with time. To prevent such an attack from being successful, keys must be changed periodically. To this end, a "session key update" procedure has been defined to support periodic key changeover. This procedure uses a master key, which is used to encrypt short-lived session keys; these in turn are used for a period of time for integrity and confidentiality services. The master key and first session key are exchanged during initial security negotiation. However, subsequent session keys must be transferred in the data channel so that the receiver may load them and start using them at the appropriate time.

The method for session key update, as described in the ATM security specification, consists of two processes: exchanging a new session key between the initiator and responder, and changing over from the old session key to the new session key. The first process is referred to as "Session Key Exchange" (SKE) and the second process is referred to as "Session Key Changeover" (SKC). The process of performing key updates is independent in each direction of data flow, for full duplex connections. It is the responsibility of the source (i.e. the encrypting side of the data confidentiality service) of each data flow to initiate the key update in its direction.

A.4.2 ATM and IP multicast security challenges

There are two important performance related considerations to be made when designing any ATM security system:

- ATM throughput: The encryption unit has to be fast and handle the full bi-directional data rates.
- Statistical multiplexing: Unique session keys are required for each VC. This requires that the cryptographic unit must be capable of changing the keys rapidly (a key agile system). Research in key agility has shown that one encryption unit for each direction can be sufficient.

Some challenges for IP multicast over ATM regarding key management are:

- Rekeying in ATM using SKC/SKE is performed in the data channel i.e. in the VC (in-band), while IP multicast systems often have a separate channel for key distribution e.g. using a different multicast address (out-of-band).
- The SKC/SKE protocols use a single ATM cell for rekeying. The size of the ATM cell restricts the use of sophisticated rekeying algorithms such as Logical Key Hierarchy, which are needed for scalability reason in large multicast groups.
- There is no true provisioning for multicast connectivity in ATM.

Annex B:
Some other satellite multicast work

B.1 SIMPLE system description

SIMPLE (http://www.simple.at) is a proxy/cache solution running both at end-user and broadcasting center side, using UDP/IP-multicast instead of TCP/IP to interactively and transparently request/reply Web based content. Upon request it replicates Web content to many sites simultaneously, hence, interactively keeping end-users caches up to date. In other words it is an Internet-recorder. Exactly like a video recorder stores interesting television programs, SIMPLE records interesting Web content on the local hard disk.

Specifically SIMPLE is a decentralized software system for distribution of information via satellite. As the name suggests, it is an easy and efficient method of multicasting multimedia Internet content, transparently delivering to and updating content on multi-user and single-user low-cost earth-stations.

SIMPLE transparently replaces TCP based point-to-point Web-browsing with UDP/IP-multicast based one-to-many Web replication and seamlessly integrates reliable multicast push services into the client cache. The quality of service (data rate and reliability) is adjusted by the operator and not but the end-user, so that SIMPLE efficiently and economically uses satellite (or any other broadcast enabled) capacity.

B.2 Spaceway

Spaceway is a next-generation satellite system that employs on-board digital processing, packet switching and spot-beam technology. It uses a mesh architecture to offer single hop connectivity between terminals on the same satellite. Bandwidth on demand is used and the system is designed to interwork seamlessly with existing terrestrial LANs and WANs.

Spaceway provides support for multicast applications using layer 2 replication in the satellite combined with spot-beam routing based on MAC addresses. In other words, the BSM satellite terminals (acting as routers) intercept and effectively spoof IGMP and PIM-SM IP signalling traffic and engage in a new satellite dependent protocol to the NCC to set up the distribution (replication) tree at L-2 as described in clause 8.6. Any terminal can source a multicast and the satellite replicates the multicast packets into all downlink beams, in which authorized terminals reside. Multicast packets have a Multicast MAC address (MGID) and all terminals that join the multicast group use this MAC address to filter the received data. Additionally, a multicast privacy key can be assigned and this is distributed to the terminals together with the associated MGID.

B.3 Hispasat

HISPASAT is mainly a space segment provider. Nevertheless HISPASAT has been actively involved in multimedia I+D projects.

In S3M project, an ACTS project in collaboration with other European companies, a multicast IP experience was performed during the years 1998 to 1999.

The multicast application involved mainly push services (fast FTP), but also video streaming (MPEG-4) and distance learning.

The main innovation was that a complete new transport was developed, RRMP (Restricted Reliable Multicast Protocol), which worked directly over IP. Basically it introduces an additional Reed-Solomon coding. This protocol increases significantly the reliability to receive error free data, clearly visible in real time video applications. The complete layers structure was RRMP over IP, MPEG-2 and DVB-S.

Nowadays, there are some users that provide IP multicast services both over Europe and America. The IP services run in a DVB-S platform, multiplexed with several video channels. Services include video streaming (UDP, MPEG 4, at 700 kb/s) and data multicast (Internet Group Management Protocol-IGMP).

B.4 SatCAST - Satellite multicast for Web applications

SatCAST is a development of software for advanced caching and replication for the WWW using multicasting, carried out by West Consulting B.V. and the University of Salzburg, in the frame of the Advanced Satellite Technology (ASTE) program of the European Space Agency.

The main components in the SatCAST proxy architecture are the Squid proxy server and the multicast agent MCast. The multicast agent is used as a peer proxy by Squid. The multicast agent uses RRMP for multicast communication with other multicast agents.

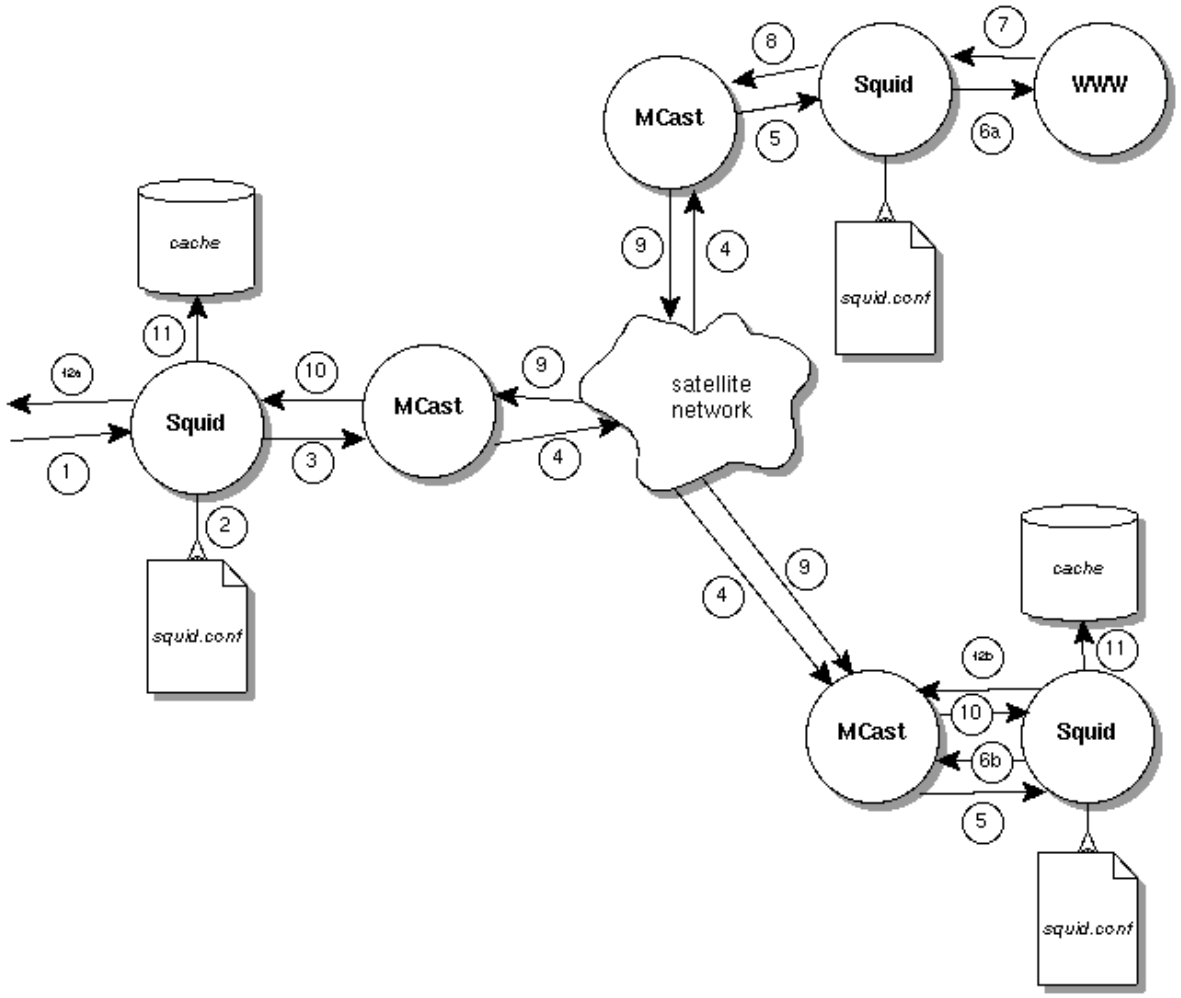

Figure B.1

The complete process for retrieving a document using Squid and MCast is shown step by step in figure B.1. Explanation of the steps in this figure:

1) A client sends an HTTP request for a URL.

2) Squid looks up in its configuration the peer that supposedly has this URL.

3) Squid forwards the HTTP request to the local multicast agent MCast.

4) MCast forwards the request by multicast to the satellite network.

5) MCast agents on other sides of the link receive the request and forward it as an HTTP request to their local Squid.

6) Squid looks up in its configuration whether the requested URL is local:

   a) and if so, sends the HTTP request to the local WWW server that holds the data.

   b) and if not, sends the HTTP request to its local MCast agent.

The next step is step 10]

7) the local WWW server sends the HTTP response to Squid.

8) Squid sends the HTTP response to its local MCast agent.

9) the MCast agent sends the response by multicast to the satellite network.

10) MCast agents on other sides of the link receive the response and forward it as an HTTP response to their local Squid.

11) Squid stores the response in the local cache.

12) Squid sends the HTTP response:

   a) to the client (if the client sent the HTTP request).

   b) to the MCast agent (if the MCast agent sent the HTTP request).

Note that this is unnecessary, but it makes the protocol consistent and it can be implemented without changes to Squid.

Clients can be users with a normal WWW browser in a network local to the Squid proxy server, or special agents like a prefetcher or push agent.

- The conclusions stated in [http://www.west.nl/whitepapers/SatCAST/TechArticle.html#Using-multicast-for-cache-replication] state that caching and replication can be very effective in reducing the perceived time of retrieval of pages from the World Wide Web and reducing the traffic from multiple requests. SatCAST software combines caching and replication with multicasting, which in turn allows to efficiently update clusters of caches. Because of the nature of multicasting, guaranteeing integrity of the data to be replicated is more complex than in unicast communications. The choice made in SatCAST has been to include a transport protocol (RRMP) running over UDP and providing the required reliability by using both FEC true provisioning for multicast connectivity in ATM.

B.5 GEOCAST study presentation

GEOCAST (Multicast over Geostationary Satellites, (http://www.geocast-satellite.com)) is a research study carried out in the area of IP multicasting over GEO satellites within the European IST program supported by the EU 5th framework program.

Different topics related to multicasting over satellite have been investigated and implemented on a real-time emulation platform:

Multicast group management:

In a DVB context, the multicast group management has been studied in order to answer to short term (static multicast) and long term objectives (dynamic multicast).

The static multicast enables to offer a selective broadcast on the air interface.

- No IP multicast protocol on the air interface.
- IGMPv2 querier in the satellite terminal as defined by the IETF RFC 2236 IGMPv2 [8].
- Definition of a Multicast Mapping Table (MMT) to inform the satellite terminal about the mapping between the Multicast IP address and the layer 2 addresses.

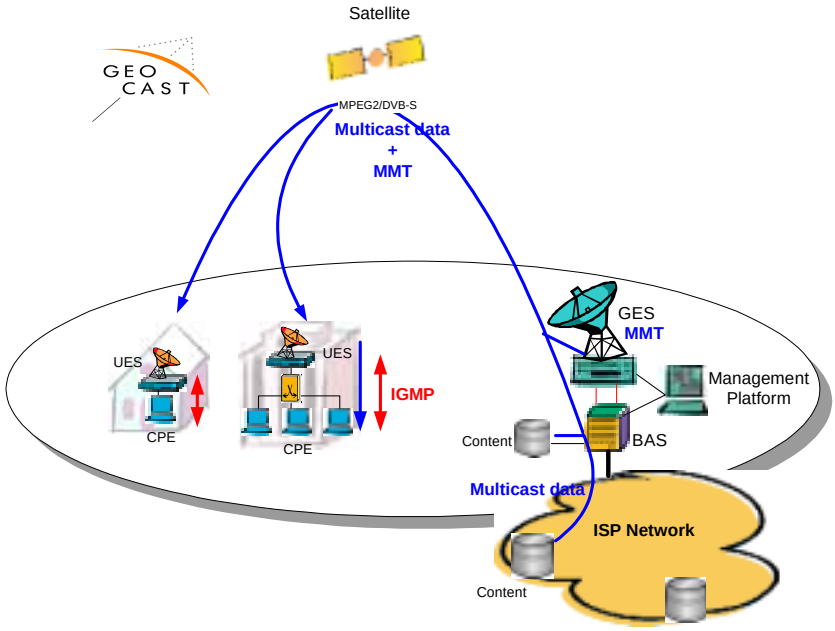

Figure B.2: GEOCAST

The dynamic multicast is much more ambitious as it enables a seamless integration of the multicast-enabled satellite network to the terrestrial networks.

- IGMP functions embedded in the satellite terminal:
  - IGMP querier as defined by the IETF RFC 2236 IGMPv2 [8].
  - IGMP proxying to relay the IGMP messages on the air interface to the first terrestrial multicast-enabled router.
- Definition of a Multicast Mapping Table (MMT) to inform the satellite terminal about the mapping between the Multicast IP address and the layer 2 addresses.

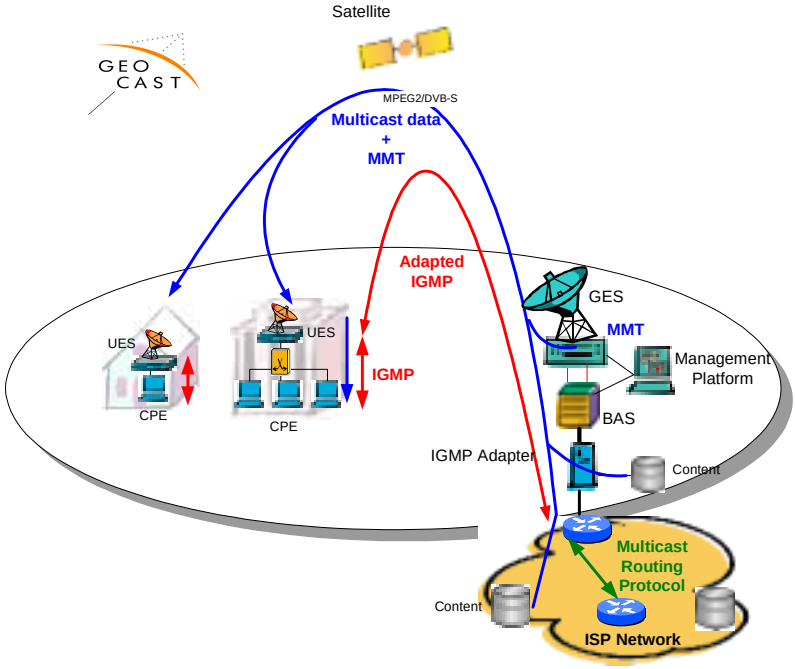

Figure B.3: GEOCAST

Because of the specificity of the satellite systems (high delay, scalability, none shared medium), some issues can be encountered: IGMP messages flooding, high latency when leaving a group.

To improve multicast services performances, some adaptations of the IGMP Querier have been specified and implemented within GEOCAST.

The main purpose is:

- to reduce the latency for stopping the traffic forwarding.
- and to reduce the number of IGMPv2 signalling messages forwarded over the air interface.

Reliable multicast transport protocol:

The multicast transport protocols have been an important research topic of the GEOCAST study. First an investigation has been performed on the behaviour of the multicast transport protocols over satellite links experiencing packet losses and high delays.

This led to the design and implementation of an optimal RMT protocol called SAT-RMTP. It is suited for deployment with large satellite terrestrial networks and compatible with the research activities proposed by the Internet Engineering Task Force (IETF) RMT working group.

Multicast security:

As IPSec is not a suitable candidate for the multicast over satellite (IKE only unicast oriented, high overhead, high latency), GEOCAST focused on the design of a layer 2 solution enabling to secure the multicast transmission on the air interface.

B.6 ICEBERGS

B.6.1 Overview

This clause (based on text input to ETSI) describes some results of a European IST project named IP ConferEncing with Broadband multimedia ovERGeostationary Satellites (ICEBERGS) about a multicast routing solution and how it can be deployed over satellite to enable multicast multiparty multimedia IP conferences. To employ the multicast over an OBP satellite is highly demanded to overcome the expensive satellite bandwidth cost.

The ICEBERGS near term solution is a protocol suite based on PIM-SM/MSDP/MBGP, which seems to be the more suitable routing architecture for the deployment of multicast routing over the ICEBERGS network. The claim is that this is because this solution can be compatible with many existing systems and supported by many ISPs while a very large scale IP conference is not that demanded in current market when considering the scalability.

This architecture consists of a set of Internet protocols herein summarized:

- The intra-domain protocol is assumed to be Protocol Independent Multicast-Sparse Mode (PIM-SM); each domain uses its own Rendezvous Point (RP) (one or more): a source located in a given domain registers to an RP of that domain.
- RPs belonging to different domains exchange information related to the existence of active sources by means of the Multicast Source Discovery Protocol (MSDP).
- An extension of the classical unicast inter-domain routing protocol, BGP is used for inter-domain routing: Multiprotocol Border Gateway Protocol, MBGP.

While MBGP is the first step toward providing inter-domain multicast, it alone is not a complete solution. MBGP is capable of determining the next hop to a host, but it has to cooperate with MSDP in order to provide multicast tree construction functions.

After studying the Models for Multi Party Conferencing in SIP (draft-ietf-sipping-conferencing-models-01), ICEBERGS defined a new conference model, Multiple Media Servers (Multiple-MCUs) Model (figure B.4), which fits both above requirements.

In this model, one or more MCUs (Multipoint Control Units) exist on the network. Terminals send multimedia streams to the MCUs by unicast, which collects the streams, manipulates them and generates multicast flows received by all terminals. This model minimizes the bandwidth in comparison to the unicast conference and simplifies the terminal requirements. The mixed satellite-terrestrial network and the relatively high satellite delay implies that it is not desirable to send unicast audio/video from one terminal to a remote MCU through a satellite and then receive the composite signal again through the satellite link.

For this reason, several MCUs are needed, at least one in each corporate/business or ISP network, so we find the "Dial-In Conference Servers" (draft-ietf-sipping-conferencing-models-01) scenario in each local network, where the Conference Server is now called MCU. In this way, an MCU acts as a normal SIP User Agent (UA): users call it, and it maintains point-to-point SIP relationships with each local-user that calls in. The MCU takes the media from the local-users who dial into the same conference, mixes them, and sends out the appropriate mixed stream to the other participant-MCUs, probably, via one satellite hop.

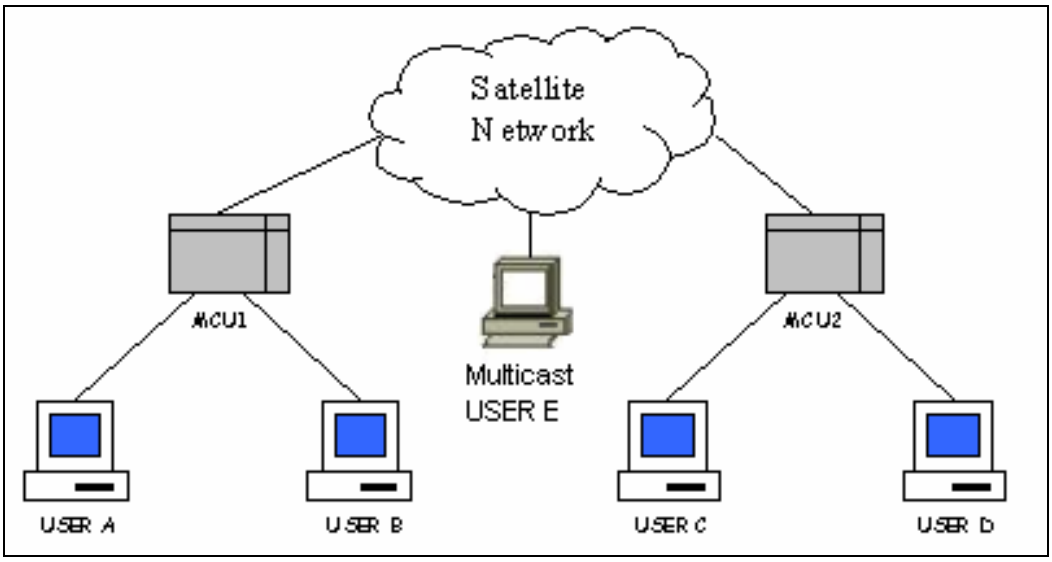

Figure B.4: MCU architecture - example

B.6.2 Intra-domain multicast routing deployment

As far as intra-domain multicast forwarding is concerned, two deployment scenarios have to be analysed in order to properly frame this protocol in the ICEBERGS context:

- PIM-SM deployment only over terrestrial networks of federated ISPs (ISPs that are not part of the satellite domain).
- PIM-SM deployment over both terrestrial networks of federated ISPs, and the satellite network.

Therefore in the case of intradomain multicast, the following scenario arises:

- PIM-SM procedures execution would result completely framed in the multicast network domain of each federated ISP; such procedures would be carried out by the Designated Router (DR) interfacing terrestrial multicast sources/receivers of a given group G towards the RPs relevant to this group located in the ISP domain; this entails that each federated ISP will autonomously manage its own RPs (no dependency on third-party RP), where group receivers will join the Rendezvous Point Tree (RPT), and group sources will join their own Shortest Path Tree (SPT).
- As far as satellite End Users (EU) are concerned, membership protocol would be executed either locally, if a multicast router is co-located with each satellite terminal (satellite enabled Multicast Router-SMR), or would be proxied at NOC if satellite terminals are not provided with an SMR. Concerning with PIM-SM, when an SMR is co-located with a satellite terminal, the relevant sessions should be accomplished between the SMR and a so called Satellite Rendezvous Point (SRP) located in the NOC, and PIM-SM sessions should be proxied at NOC as well (brown curves in figure B.5).

In particular, in the context of the ISPk network domain, the PIM-SM messages are carried out between each DR and the RP of the multicast group, which has been considered co-located with an SMR. In the frame of the multicast routing, we have only to take into account that PIM-SM sessions relevant to unicast EUs. This will take place between the Multicast Router (MR) located at ISPk Unicast-Multicast domain boundary and the RP of the multicast group. Finally, as far as satellite EUs are concerned, we have considered satellite EUs co-located with an SMR, therefore IGMP sessions may be locally carried out, while PIM-SM sessions are performed between each SMR and the SRP.

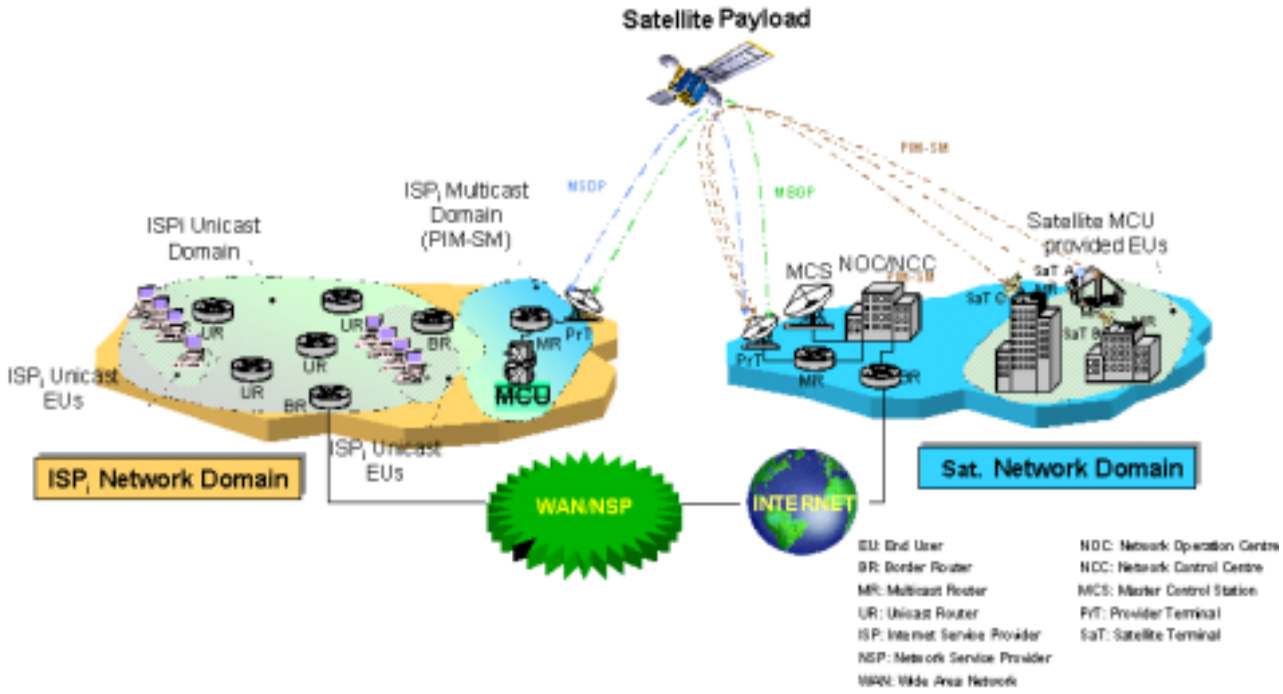

Figure B.5: ICEBERGS multicast and unicast hybrid routing architecture

Annex C:
Some useful web links

- http://www.multicasttech.com/faq/
- http://www.zvon.org/
- http://www.uoregon.edu/~llynch/calendar/
- http://www.dvb.org
- http://www.itu.int
- http://www.ietf.org

Some sites with regular active multicast content:

- UofO http://videolab.uoregon.edu/
- Access Grid http://www-fp.mcs.anl.gov/fl/accessgrid/
- ICAIR C-Span http://cspan.icair.org/
- On-the-I http://www.on-the-i.com/
- Yahoo http://www.broadcast.com/broadband/
- NASA http://www.nasa.gov/
- Berkeley http://media2.bmrc.berkeley.edu/bibs/archive.cfm
- UCSB http://imj.ucsb.edu/
- and the listing found in ftp://limestone.uoregon.edu/pub/multicast/mice/sdr/

Annex D:
Multicast related RFCs

IETF RFCs that refer to multicast, from http://www.zvon.org/:

IETF RFC 1949: "Scalable Multicast Key Distribution".

IETF RFC 3170: "IP Multicast Applications: Challenges and Solutions".

IETF RFC 2907: "MADCAP Multicast Scope Nesting State Option".

IETF RFC 2357: "IETF Criteria for Evaluating Reliable Multicast Transport and Application Protocols".

IETF RFC 2588: "IP Multicast and Firewalls".

IETF RFC 2365: "Administratively Scoped IP Multicast".

IETF RFC 1458: "Requirements for Multicast Protocols".

IETF RFC 1584: "Multicast Extensions to OSPF".

IETF RFC 2490: "A Simulation Model for IP Multicast with RSVP".

IETF RFC 3353: "Overview of IP Multicast in a Multi-Protocol Label Switching (MPLS) Environment".

IETF RFC 2375: "IPv6 Multicast Address Assignments".

IETF RFC 3306: "Unicast-Prefix-based IPv6 Multicast Addresses".

IETF RFC 2502: "Limitations of Internet Protocol Suite for Distributed Simulation the Large Multicast Environment".

IETF RFC 2201: "Core Based Trees (CBT) Multicast Routing Architecture".

IETF RFC 2366: "Definitions of Managed Objects for Multicast over UNI 3.0/3.1 based ATM Networks (Obsoleted by RFC 2417)".

IETF RFC 2417: "Definitions of Managed Objects for Multicast over UNI 3.0/3.1 based ATM Networks".

IETF RFC 2627: "Key Management for Multicast: Issues and Architectures".

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

IETF RFC 2102: "Multicast Support for Nimrod: Requirements and Solution Approaches".

IETF RFC 1469: "IP Multicast over Token-Ring Local Area Networks".

IETF RFC 2090: "TFTP Multicast Option".

IETF RFC 2715: "Interoperability Rules for Multicast Routing Protocols".

IETF RFC 2730: "Multicast Address Dynamic Client Allocation Protocol (MADCAP)".

IETF RFC 2771: "An Abstract API for Multicast Address Allocation".

IETF RFC 3307: "Allocation Guidelines for IPv6 Multicast Addresses".

IETF RFC 0966: "Host groups: A multicast extension to the Internet Protocol (Obsoleted by RFC 0988 -> RFC 1054 -> RFC 1112; -> RFC 1112)".

IETF RFC 2710: "Multicast Listener Discovery (MLD) for IPv6".

IETF RFC 1301: "Multicast Transport Protocol".

IETF RFC 2337: "Intra-LIS IP multicast among routers over ATM using Sparse Mode PIM".

IETF RFC 2991: "Multipath Issues in Unicast and Multicast Next-Hop Selection".

IETF RFC 2432: "Terminology for IP Multicast Benchmarking".

IETF RFC 2674: "Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions".

IETF RFC 2908: "The Internet Multicast Address Allocation Architecture".

IETF RFC 2932: "IPv4 Multicast Routing MIB".

IETF RFC 2934: "Protocol Independent Multicast MIB for IPv4".

IETF RFC 2909: "The Multicast Address-Set Claim (MASC) Protocol".

IETF RFC 1075: "Distance Vector Multicast Routing Protocol".

IETF RFC 3048: "Reliable Multicast Transport Building Blocks for One-to-Many Bulk-Data Transfer".

IETF RFC 3269: "Author Guidelines for Reliable Multicast Transport (RMT) Building Blocks and Protocol Instantiation documents".

IETF RFC 2776: "Multicast-Scope Zone Announcement Protocol (MZAP)".

IETF RFC 2149: "Multicast Server Architectures for MARS-based ATM multicasting".

IETF RFC 2887: "The Reliable Multicast Design Space for Bulk Data Transfer".

IETF RFC 3019: "IP Version 6 Management Information Base for The Multicast Listener Discovery Protocol".

IETF RFC 2729: "Taxonomy of Communication Requirements for Large-scale Multicast Applications".

IETF RFC 2189: "Core Based Trees (CBT version 2) Multicast Routing -- Protocol Specification".

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

IETF RFC 3082: "Notification and Subscription for SLP".

IETF RFC 1340: "Assigned Numbers (Obsoleted by RFC 1700 -> RFC 3232)".

IETF RFC 1060: "Assigned numbers (Obsoleted by RFC 1340 -> RFC 1700 -> RFC 3232) (Updated by RFC 1349)".

IETF RFC 1054: "Host extensions for IP multicasting (Obsoleted by RFC 1112)".

IETF RFC 1112: "Host extensions for IP multicasting (Updated by RFC 2236)".

IETF RFC 2529: "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels".

IETF RFC 2226: "IP Broadcast over ATM Networks".

IETF RFC 1700: "Assigned Numbers (Obsoleted by RFC 3232)".

IETF RFC 0988: "Host extensions for IP multicasting (Obsoleted by RFC 1054 -> RFC 1112, RFC 1112)".

IETF RFC 2166: "APPN Implementer's Workshop Closed Pages Document DLSw v2.0 Enhancements".

IETF RFC 1768: "Host Group Extensions for CLNP Multicasting".

IETF RFC 1812: "Requirements for IP Version 4 Routers (Updated by RFC 2644)".

IETF RFC 2174: "A MAPOS version 1 Extension - Switch-Switch Protocol".

IETF RFC 1716: "Towards Requirements for IP Routers (Obsoleted by RFC 1812)".

IETF RFC 2094: "Group Key Management Protocol (GKMP) Architecture".

IETF RFC 0992: "On communication support for fault tolerant process groups".

IETF RFC 2386: "A Framework for QoS-based Routing in the Internet".

IETF RFC 2380: "RSVP over ATM Implementation Requirements".

IETF RFC 2734: "IPv4 over IEEE 1394".

IETF RFC 3111: "Service Location Protocol Modifications for IPv6".

IETF RFC 2998: "A Framework for Integrated Services Operation over Diffserv Networks".

IETF RFC 2834: "ARP and IP Broadcast over HIPPI-800".

IETF RFC 2835: "IP and ARP over HIPPI-6400 (GSN)".

IETF RFC 2746: "RSVP Operation Over IP Tunnels".

IETF RFC 2470: "Transmission of IPv6 Packets over Token Ring Networks".

IETF RFC 2475: "An Architecture for Differentiated Service (Updated by RFC 3260)".

IETF RFC 2497: "Transmission of IPv6 Packets over ARCnet Networks".

IETF RFC 3251: "Electricity over IP".

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

IETF RFC 2467: "Transmission of IPv6 Packets over FDDI Networks".

IETF RFC 2019: "Transmission of IPv6 Packets Over FDDI (Obsoleted by RFC 2467)".

IETF RFC 2121: "Issues affecting MARS Cluster Size".

IETF RFC 3091: "Pi Digit Generation Protocol".

IETF RFC 3146: "Transmission of IPv6 Packets over IEEE 1394 Networks".

IETF RFC 1088: "Standard for the transmission of IP datagrams over NetBIOS networks".

IETF RFC 1577: "Classical IP and ARP over ATM (Obsoleted by RFC 2225)".

IETF RFC 1671: "IPng White Paper on Transition and Other Considerations".

IETF RFC 1677: "Tactical Radio Frequency Communication Requirements for IPng".

IETF RFC 1070: "Use of the Internet as a subnetwork for experimentation with the OSI network layer".

IETF RFC 3056: "Connection of IPv6 Domains via IPv4 Clouds".

IETF RFC 2590: "Transmission of IPv6 Packets over Frame Relay Networks Specification".

IETF RFC 2768: "Network Policy and Services: A Report of a Workshop on Middleware".

IETF RFC 1209: "Transmission of IP datagrams over the SMDS Service".

IETF RFC 2175: "MAPOS 16 - Multiple Access Protocol over SONET/SDH with 16 Bit Addressing".

IETF RFC 3069: "VLAN Aggregation for Efficient IP Address Allocation".

IETF RFC 1188: "Proposed Standard for the Transmission of IP Datagrams over FDDI Networks".

IETF RFC 2225: "Classical IP and ARP over ATM".

IETF RFC 1707: "CATNIP: Common Architecture for the Internet".

IETF RFC 2187: "Application of Internet Cache Protocol (ICP), version 2".

IETF RFC 1390: "Transmission of IP and ARP over FDDI Networks".

IETF RFC 1546: "Host Anycasting Service".

IETF RFC 1972: "A Method for the Transmission of IPv6 Packets over Ethernet Networks (Obsoleted by RFC 2464)".

IETF RFC 2022: "Support for Multicast over UNI 3.0/3.1 based ATM Networks".

IETF RFC 2491: "IPv6 over Non-Broadcast Multiple Access (NBMA) networks".

IETF RFC 2362: "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification".

IETF RFC 2117: "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification (Obsoleted by RFC 2362)".

IETF RFC 2165: "Service Location Protocol (Updated by RFC 2608, RFC 2609)".

IETF RFC 2543: "SIP: Session Initiation Protocol (Obsoleted by RFC 3261, RFC 3262, RFC 3263, RFC 3264, RFC 3265)".

IETF RFC 2814: "SBM (Subnet Bandwidth Manager): A Protocol for RSVP-based Admission Control over IEEE 802-style networks".

IETF RFC 3264: "An Offer/Answer Model with Session Description Protocol (SDP)".

IETF RFC 1825: "Security Architecture for the Internet Protocol (Obsoleted by RFC 2401)".

IETF RFC 3031: "Multiprotocol Label Switching Architecture".

IETF RFC 1621: "Pip Near-term Architecture".

IETF RFC 2009: "GPS-Based Addressing and Routing".

IETF RFC 2382: "A Framework for Integrated Services and RSVP over ATM".

IETF RFC 2902: "Overview of the 1998 IAB Routing Workshop".

IETF RFC 2373: "IP Version 6 Addressing Architecture".

IETF RFC 2764: "A Framework for IP Based Virtual Private Networks".

IETF RFC 3102: "Realm Specific IP: Framework".

IETF RFC 3259: "A Message Bus for Local Coordination".

IETF RFC 3175: "Aggregation of RSVP for IPv4 and IPv6 Reservations".

IETF RFC 1884: "IP Version 6 Addressing Architecture (Obsoleted by RFC 2373)".

IETF RFC 1726: "Technical Criteria for Choosing IP The Next Generation (IPng)".

IETF RFC 3379: "Delegated Path Validation and Delegated Path Discovery Protocol Requirements".

IETF RFC 2326: "Real Time Streaming Protocol (RTSP)".

IETF RFC 2327: "SDP: Session Description Protocol (Updated by RFC 3266)".

IETF RFC 1045: "VMTP: Versatile Message Transaction Protocol".

IETF RFC 3154: "Requirements and Functional Architecture for an IP Host Alerting Protocol".

IETF RFC 1889: "RTP: A Transport Protocol for Real-Time Applications".

IETF RFC 2340: "Nortel's Virtual Network Switching (VNS) Overview".

IETF RFC 2344: "Reverse Tunneling for Mobile IP (Obsoleted by RFC 3024)".

IETF RFC 2002: "IP Mobility Support (Obsoleted by RFC 3220 -> RFC 3344) (Updated by RFC 2290)".

IETF RFC 2133: "Basic Socket Interface Extensions for IPv6 (Obsoleted by RFC 2553)".

IETF RFC 1329: "Thoughts on Address Resolution for Dual MAC FDDI Networks".

IETF RFC 1679: "HPN Working Group Input to the IPng Requirements Solicitation".

IETF RFC 3049: "TN3270E Service Location and Session Balancing".

IETF RFC 2961: "RSVP Refresh Overhead Reduction Extensions".

IETF RFC 1788: "ICMP Domain Name Messages".

IETF RFC 2749: "COPS usage for RSVP".

IETF RFC 1636: "Report of IAB Workshop on Security in the Internet Architecture", February 8-10, 1994.

IETF RFC 2105: "Cisco Systems' Tag Switching Architecture Overview".

IETF RFC 2101: "IPv4 Address Behaviour Today".

IETF RFC 1682: "IPng BSD Host Implementation Analysis".

IETF RFC 2114: "Data Link Switching Client Access Protocol".

IETF RFC 2546: "6Bone Routing Practice (Obsoleted by RFC 2772)".

IETF RFC 1201: "Transmitting IP traffic over ARCNET networks".

IETF RFC 3024: "Reverse Tunneling for Mobile IP, revised".

IETF RFC 2608: "Service Location Protocol, Version 2 (Updated by RFC 3224)".

IETF RFC 2176: "IPv4 over MAPOS Version 1".

IETF RFC 2772: "6Bone Backbone Routing Guidelines (Updated by RFC 3152)".

IETF RFC 2553: "Basic Socket Interface Extensions for IPv6 (Updated by RFC 3152)".

IETF RFC 1454: "Comparison of Proposals for Next Version of IP".

IETF RFC 1585: "MOSPF: Analysis and Experience".

IETF RFC 1721: "RIP Version 2 Protocol Analysis".

IETF RFC 1723: "RIP Version 2 - Carrying Additional Information (Obsoleted by RFC 2453)".

IETF RFC 1722: "RIP Version 2 Protocol Applicability Statement".

IETF RFC 1821: "Integration of Real-time Services in an IP-ATM Network Architecture".

IETF RFC 2093: "Group Key Management Protocol (GKMP) Specification".

IETF RFC 2097: "The PPP NetBIOS Frames Control Protocol (NBFCP)".

IETF RFC 3344: "IP Mobility Support for IPv4".

IETF RFC 1009: "Requirements for Internet gateways (Obsoleted by RFC 1812)".

IETF RFC 0824: "CRONUS Virtual Local Network".

IETF RFC 2408: "Internet Security Association and Key Management Protocol (ISAKMP)".

IETF RFC 2401: "Security Architecture for the Internet Protocol (Updated by RFC 3168)".

IETF RFC 1946: "Native ATM Support for ST2+".

IETF RFC 3220: "IP Mobility Support for IPv4 (Obsoleted by RFC 3344)".

IETF RFC 1387: "RIP Version 2 Protocol Analysis (Obsoleted by RFC 1721)".

IETF RFC 1388: "RIP Version 2 Carrying Additional Information (Obsoleted by RFC 1723 -> RFC 2453)".

IETF RFC 3234: "Middleboxes: Taxonomy and Issues".

IETF RFC 2453: "RIP Version 2".

IETF RFC 1754: "IP over ATM Working Group's Recommendations for the ATM Forum's Multiprotocol BOF Version 1".

IETF RFC 2780: "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers".

IETF RFC 1573: "Evolution of the Interfaces Group of MIB-II (Obsoleted by RFC 2233 -> RFC 2863)".

IETF RFC 2492: "IPv6 over ATM Networks".

IETF RFC 1710: "Simple Internet Protocol Plus White Paper".

IETF RFC 1077: "Critical issues in high bandwidth networking".

IETF RFC 3235: "Network Address Translator (NAT)-Friendly Application Design Guidelines".

IETF RFC 2383: "ST2+ over ATM Protocol Specification - UNI 3.1 Version".

IETF RFC 2233: "The Interfaces Group MIB using SMIv2 (Obsoleted by RFC 2863)".

IETF RFC 2150: "Humanities and Arts: Sharing Center Stage on the Internet".

IETF RFC 1001: "Protocol standard for a NetBIOS service on a TCP/UDP transport: Concepts and methods".

IETF RFC 1122: "Requirements for Internet Hosts - Communication Layers (Updated by RFC 1349)".

IETF RFC 1819: "Internet Stream Protocol Version 2 (ST2) Protocol Specification - Version ST2+".

IETF RFC 1190: "Experimental Internet Stream Protocol: Version 2 (ST-II) (Obsoleted by RFC 1819)".

IETF RFC 2863: "The Interfaces Group MIB".

History

Document history		
V1.1.1	April 2003	Publication