ETSI TR 101 985 V1.1.1 (2002-11)

Technical Report

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

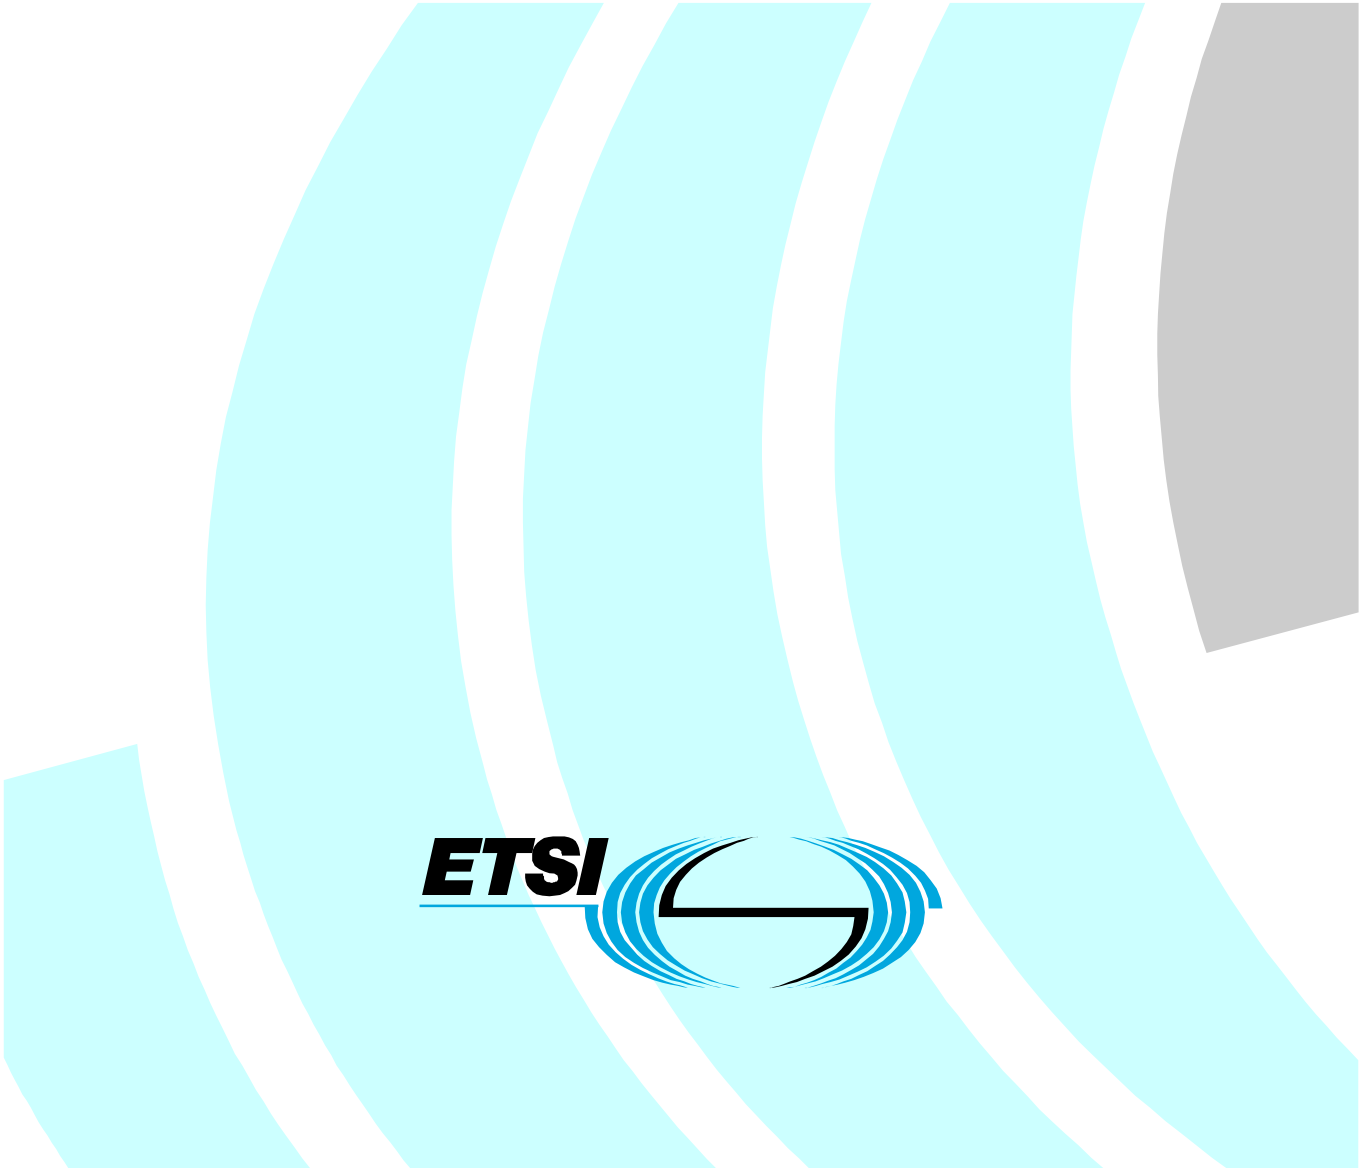

Reference
DTR/SES-00064

Keywords
architecture, 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

Important notice

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

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

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

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

Copyright Notification

Contents

Intellectual Property Rights

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

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

Foreword

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

1 Scope

The scope of the present document is the review of study topics with respect to the provision of IP-based services via Broadband Satellite Multimedia (BSM) networks. The present document builds on the following two earlier reports:

- TR 101 374-1 [1], "Survey on Standardization Objectives for Broadband Satellite Multimedia".
- TR 101 374-2 [2], "Standardization Objectives for Broadband Satellite Multimedia: The Standardization Scenario".

The present document also builds on the general BSM Services and Architectures defined in:

- TR 101 984 [3], "Services and Architectures".

The present document discusses the standardization approach, relevant issues and provides a reference framework for further work that should be undertaken within ETSI with regard to the use of Broadcast Satellite Multimedia (BSM) networks to transport IP-structured traffic. The report is focussed on the following aspects off IP over satellite:

- BSM systems based on GeoSynchronous Orbit (GSO) satellites;
- integration of BSM services with IP-based services;
- integration of BSM satellite networks with terrestrial networks;
- integration of BSM broadcast and multicast services with IP-based services, including IP multicast services.

The structure defined in the present document is intended to provide a framework for a series of detailed studies on specific issues which are expected to identify existing standards and to identify areas where new standards are required, in order to achieve these objectives. The objectives of BSM standardization are:

- to enable users to access a wide range of telecommunications services, with particular emphasis on multi-media services and high data rates;
- to provide an efficient means of using network resources (particularly radio spectrum);
- to enable the benefits of satellite technology to be made available to a wide range of users;
- to facilitate the provision of a high quality of service for transporting IP traffic over BSM networks;
- to facilitate the provision of low cost terminals.

2 References

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

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

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

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

[4] ISO/IEC 3309 (1993): "Information Technology - Telecommunications and information exchange between systems - High level data link control (HDLC) procedures - Frame structure".

[5] ISO/IEC 7498: "Information Technology - Open Systems Interconnection - Basic Reference Model".

[6] ETSI TS 123 107: "Universal Mobile Telecommunications System (UMTS);QoS Concept and Architecture(Release 1999)".

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

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

[9] IETF RFC 1633: "Integrated Services in the Internet Architecture: an Overview".

[10] IETF RFC 2475: "An Architecture for Differentiated Service".

[11] IETF RFC 2990: "Next steps for the IP QoS Architecture".

[12] IETF RFC 2205: "Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification".

[13] IETF RFC 2210: "The Use of RSVP with IETF Integrated Services".

[14] IETF RFC 2211: "Specification of the Controlled-Load Network Element Service".

[15] IETF RFC 2212: "Specification of Guaranteed Quality of Service".

[16] IETF RFC 2474: "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers".

[17] IETF RFC 2597: "Assured Forwarding PHB Group".

[18] IETF RFC 2598: "An Expedited Forwarding PHB".

[19] IETF RFC 2208: "Resource ReSerVation Protocol (RSVP) - Version 1 Applicability Statement Some Guidelines on Deployment".

[20] IETF RFC 3135: "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations".

[21] IETF RFC 826: "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware".

[22] IETF RFC 2461: "Neighbour Discovery for IP Version 6 (IPv6)".

[23] IETF RFC 2960: "Stream Control Transmission Protocol".

[24] IETF RFC 2616: "Hypertext Transfer Protocol - HTTP/1.1".

[25] IETF RFC 1191: "Path MTU discovery".

[26] IETF RFC 3168: "The Addition of Explicit Congestion Notification (ECN) to IP".

[27] IETF RFC 2402: "IP Authentication Header".

[28] IETF RFC 2406: "IP Encapsulating Security Payload (ESP)".

[29] IETF RFC 2409: "The Internet Key Exchange (IKE)".

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

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

[32] IETF RFC 2488: "Enhancing TCP Over Satellite Channels using Standard Mechanisms".

[33] draft-ietf-pilc-link-design-12 (2002): "Advice for Internet Subnetwork Designers", Phil Karn.

[34] IETF Protocol Independent Multicast (pim) working group; www.ietf.org.

[35] IETF Multicast Security (msec) working group; www.ietf.org.

[36] IETF Multicast & Anycast Group Membership (magma) working group; www.ietf.org.

[37] IETF Performance Implications of Link Characteristics (pilc) working group, www.ietf.org.

[38] ITU-R Recommendation P.618: "Propagation data and prediction methods required for the design of Earth-space telecommunication systems".

[39] ITU-R Recommendation P.837: "Characteristics of precipitation for propagation modelling".

[40] ITU-R Recommendation P.838: "Specific attenuation model for rain for use in prediction methods".

[41] ITU-R Recommendation P.839: "Rain height model for prediction methods".

[42] ITU-T Recommendation Y.1541: "Network performance objectives for IP-Based services".

[43] ITU-T Recommendation Y.1540: "Internet protocol data communication service - IP packet transfer and availability performance parameters".

[44] ITU-T Draft Recommendation Y.iptc: "Traffic Control and Congestion Control in IP based Networks".

[45] Space Communications Protocol Standards (SCPS) website; www.scps.org.

[46] ISO 15893 (2000): "Space data and information transfer systems - Protocol specification for space communications - Transport protocol".

[47] Military Standard (MIL-STD-2045-44000): "DoD Interface Standard: Transport Protocol".

3 Definitions and abbreviations

3.1 Definitions

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

service attribute: specified characteristic of a telecommunication service

NOTE: The value(s) assigned to one or more service attributes may be used to distinguish that telecommunication service from others.

service category or service class: service offered to the users described by a set of performance parameters and their specified values, limits or ranges

NOTE: The set of parameters provides a comprehensive description of the service capability.

service component: single type of telecommunication service

NOTE: Service components are divided into speech, audio, video and data:

Speech: voice telecommunication.

Audio: telecommunication of sound in general.

Video: telecommunication of full motion pictures, and of stills.

Data: telecommunication of information-files (text, graphics, etc).

MultiMedia (MM): a combination of two or more of the above components (speech, audio, video, data), with a temporal relationship (e.g. synchronization) between at least two components.

telecommunication service: service offered by a network operator or service provider to its customers in order to satisfy a specific telecommunication requirement

NOTE: Telecommunication services are divided into two broad families: bearer services and teleservices:

bearer service: is a type of telecommunication service that provides the capability of transmission of signals between access points.

teleservice: is a type of telecommunication service that provides the complete capability, including terminal equipment functions, for communication between users according to standardized protocols and transmission capabilities established by agreement between operators.

3.2 Abbreviations

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

ACK	acknowledgement
AF	Assured Forwarding
AH	Authentication Header
AR	Address Resolution
ARP	Address Resolution Protocol
ATM	Asynchronous Transfer Mode
BE	Best-Effort
BER	Bit Error Ratio
BGP	Border Gateway Protocol
BSM	Broadband Satellite Multimedia
CA	Conditional Access
CHAP	Challenge Handshake Authentication Protocol
CLS	Controlled Load Service
CPE	Customer Premises Equipment
CRC	Cyclic Redundancy Check
DBW	Dedicated BandWidth
DCCP	Datagram Congestion Control Protocol
Diffserv	Differentiated Services
DSCP	Diffserv Codepoint
DVB	Digital Video Broadcasting
DVB-RCS	DVB-Return Channel by Satellite
ECN	Explicit Cngestion Notification
EF	Expedited Forwarding
EGP	Exterior Gateway Protocol
ES	End System
ESP	Encapsulated Security Payload
FTP	File Transfer Protocol
GS	Guaranteed Service
GSO	Geo-Stationary Orbit
HTTP	HyperText Transfer Protocol
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IKE	Internet Key Exchange
Intserv	Integrated Services
IP	Internet Protocol
IPR	Intellectual Property Rights
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISO	International Organization for Standardization
ISP	Internet Service Provider
ITU	International Telecommunication Union
ITU-R	ITU Radiocommunication sector
ITU-T	ITU Telecommunication Standardization sector
LAN	Local Area Network
MAC	Medium Access Control

MF-TDMA	Multi-Frequency TDMA
MLD	Multicast Listener Discovery
MMT	Multicast Mapping Table
MPEG	Moving Picture Expert Group
MSL	Maximum Segment Lifetime
MTU	Maximum Transmission Unit
NAT	Network Address Translation
NCC	Network Control Centre
ND	Neighbour Discovery
OBP	On Board Processing
OSPF	Open Shortest Path First
PAP	Password Authentication Protocol
PDU	Protocol Data Unit
PEP	Performance Enhancing Proxy
PGP	Pretty Good Privacy
PHB	Per-Hop Behaviour
PID	DVB Packet IDentifier
PIM-SM	Protocol Independent Multicast - Sparse Mode
QoS	Quality of Service
RFC	IETF Request For Comments
RIP	Routing Information Protocol
RP	Rendez-vous Point
RSVP	Reservation Protocol
SA	Security Association
SAF	Satellite Access Function
SAP	Session Announcement Protocol
SBW	Statistical BandWidth
SCPS	Space Communications Protocol Standards
SCPS-TP	SCPS Transport Protocol
SCTP	Stream Control Transmission Protocol
SDAF	Satellite Dependent Adaptation Functions
SDU	Service Data Unit
SGF	Satellite Gateway Function
SIAF	Satellite Independent Adaptation Functions
SI-SAP	Satellite Independent-Service Access Point
SLC	Satellite Link Control
SOHO	Small Office - Home Office
SPHY	Satellite Physical
SPI	Security Parameter Index
SSL	Secure Socket Layer
ST	Satellite Terminal
STL	Satellite Transport Lane
TCP	Transmission Control Protocol
TCS	Traffic Conditioning Specification
TDMA	Time Division Multiple Access
TIPHON	Telecommunications and Internet Protocol Harmonization Over Networks
TLS	Transport Layer Security
TTL	Time To Live
UDP	User Datagram Protocol
UNI	User Network Interface
UTL	Upper Transport Layer
VP	Virtual Path
WG	Working Group

4 Overview

The present document divides the discussion of IP over satellite into a series of inter-related study areas. These areas are based on a combination of the ITU IP-project organization and the 3GPP document organization. Each study area is discussed in more detail in a separate clause. Table 4.1 lists all the clauses together with a brief description of the contents of each clause.

Table 4.1: Grouping of IP study areas

Study Area	Clause	Description
Reference Architectures	5	Architectures and associated definitions for the specific cases of IP interworking. Based on the Services and Architecture models [3].
Bearer services for transport for IP structured signals	6	Basic requirements for transport of IP packets
Performance and Availability	7	Top level generic performance requirements: delay class, availability class, etc.
Quality of Service (QoS)	8	QOS for a BSM network, using an architecture based on TS 123 107 [6]
Routing and Addressing	9	Routing issues, address resolution, etc.
Multicast and Broadcast	10	Multicast and Broadcast issues, with particular reference to efficient handling of IP multicast and IP streaming
Security	11	Security threats and requirements Satellite independent mechanisms (incl. IPSec) plus satellite dependent mechanisms.
Performance Enhancing Proxies (PEPs)	12	A review of PEPs, based on IETF RFC 3135 [20]

Clause 5 defines the basic set of IP interworking architectures that are used in the rest of the report. This clause builds on the BSM services and architectures report [3] by mapping the generic BSM services and models into IP specific services and models.

The rest of the present document (clauses 6 to 12) discusses a series of separate IP-related issues - each clause deals with one study area. It is intended that these clauses will provide a starting point for more detailed reports on each of these issues.

5 Reference architectures

5.1 IP scenarios

The IP scenarios are defined in the Services and Architectures report [3].

The IP scenarios are grouped into three different use cases:

- Access Network; including point-to-point, multicast and broadcast services.
- Content Distribution to the Edge; including point-to-point and multicast services.
- Core Network; including point-to-point services only.

A BSM network can support all three scenarios. However, the present document will give priority to issues related to the first two scenarios, namely Access Network scenarios and Content Distribution to the Edge.

5.2 Reference architectures

5.2.1 Definitions

A BSM network may support either a mesh or star topology as defined in the Services and Architectures [3]:

- A star network topology is defined by the star arrangement of links between the Hub station (or Gateway) and multiple Remote stations. A Remote station can only establish a direct link with the Hub station and cannot establish a direct link to another Remote station.
- A mesh network is defined by the mesh arrangement of links between the stations, where any station can link directly to any other station. The star topology can be considered as one special case of the mesh topology.

NOTE: A star topology can be used to provide mesh connectivity by establishing an indirect link between Remote stations via the Hub station.

A BSM network may use either a non-regenerative or a regenerative satellite architecture:

- A non-regenerative architecture refers to a single architecture, commonly called a "bent-pipe architecture". This architecture does not terminate any layers of the air interface protocol stack in the satellite - the satellite simply transfers the signals from the user links to the feeder links transparently.
- A regenerative architecture is the range of other architectures that provide additional functionality in the satellite. In these architectures, the satellite functions terminate one or more layers of the air interface protocol stack in the satellite.

5.2.2 Reference model

The reference models are defined in the BSM services and architectures report [3].

5.3 Protocol architecture

5.3.1 Definitions

The protocol architecture is defined in the Services and Architectures report [3] and is reproduced in figure 5.3.1.

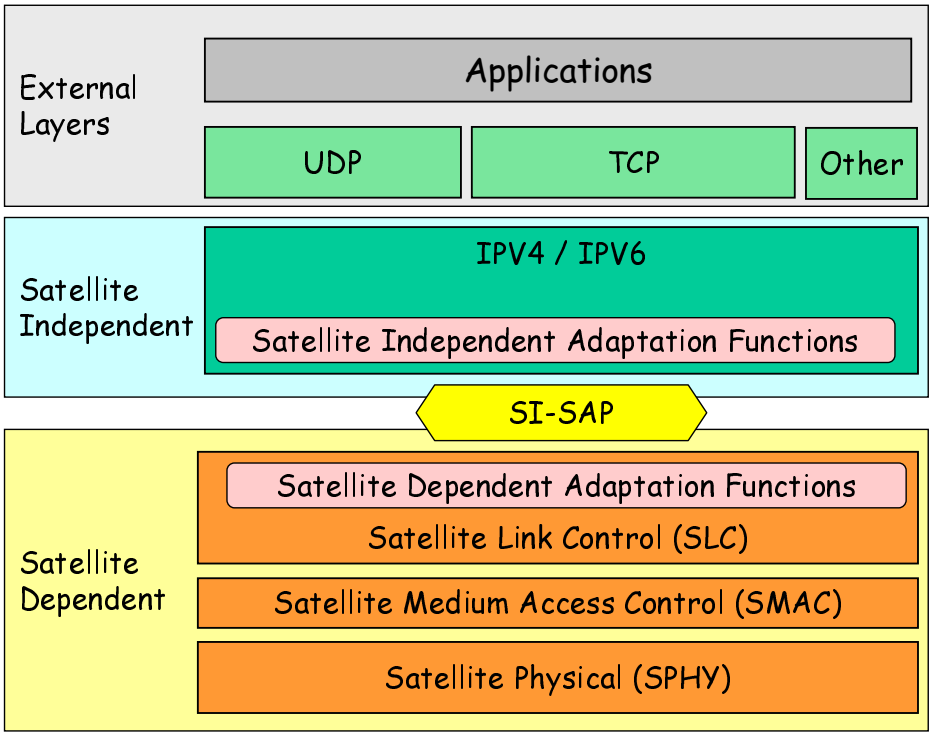

Figure 5.3.1: Protocol architecture

A satellite independent service access point (SI-SAP) is defined for the air interface to separate the satellite independent upper layers from the satellite dependent lower layers.

5.3.2 SI-SAP reference model

Figure 5.3.2 shows a more detailed reference model for the protocol architecture. The protocol stack is divided into the lower, satellite dependent layers and the upper satellite independent layers. These two parts are connected via the Satellite Independent interface (SI-SAP). The SI-SAP is logically divided into three functional SAPs as illustrated in figure 5.3.2:

- The SI-U-SAP for the U-plane services
- The SI-C-SAP for the C-plane services
- The SI-M-SAP for the M-plane services

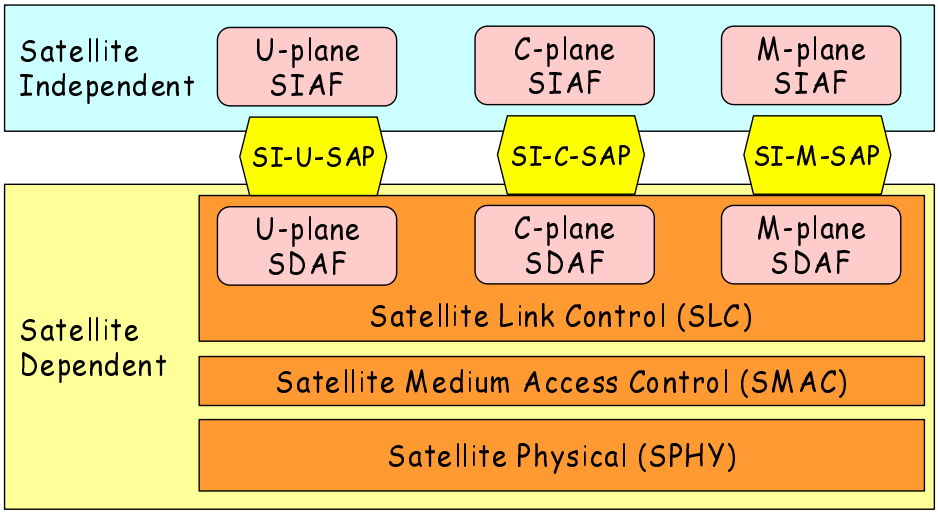

Figure 5.3.2: SI-SAP reference model

A pair of adaptation functions (an upper Satellite Independent Adaptation Function (SIAF) and a lower Satellite Dependent Adaptation Function (SDAF)) is associated with each logical SAP.

5.3.3 Interworking models for satellite subnetworks

5.3.3.1 General

A BSM subnetwork can interwork with the external IP subnetworks at different levels in the IP protocol stack. The present document defines the following different cases:

- Bridge interworking function (interworking below the IP layer);
- IP interworking function (interworking at the IP layer);
- Gateway function (interworking above the IP layer).

These interworking options are functional definitions and a given ST may provide different types of interworking for different applications or for different users.

NOTE: Other cases of interworking are also possible.

5.3.3.2 Bridge interworking functions

A bridge interworking function is defined as operating below the IP level of the protocol stack.

In the case of terrestrial subnetworks, a bridge interworking function is used to interconnect LAN segments. A bridge differs from a repeater by providing some additional functions. For example, a bridge should only forward valid (i.e. error checked) frames. A bridge may also provide filtering of frames based on the MAC addresses.

In the case of a satellite subnetwork, the bridge interworking functions include frame interworking: the external subnetwork frame format must be interworked into the satellite subnetwork frame format as illustrated in figure 5.3.3.2. The external frames are terminated in the ST and the IP payload is transferred into the different satellite frames.

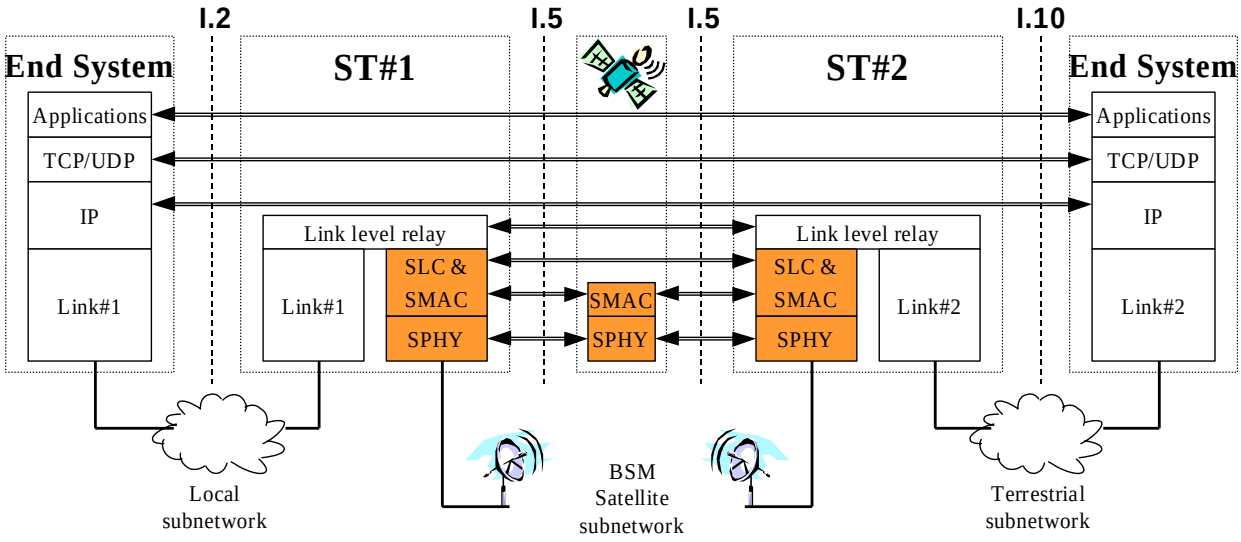

Figure 5.3.3.2: Bridge interworking functions

5.3.3.3 IP interworking function

An IP interworking function is defined as operating at the IP level of the protocol stack.

In the case of terrestrial subnetwork, the IP interworking function are typically provided by an IP router. Many different types of router are possible, with different levels of IP interworking functionality, depending on the position of the router in the overall hierarchy of IP networks.

In the case of a satellite subnetwork, the routing functions are broadly similar to the terrestrial router functions. However, many of the terrestrial routing protocols are not suitable for direct use over a satellite subnetwork and the satellite subnetwork is therefore required to use different protocols internally and the IP interworking functions illustrated in figure 5.3.3.3 are used to interwork between these different external and internal protocols.

The routed IP packets can be interworking into the same satellite subnetwork frame format as used for the bridge function.

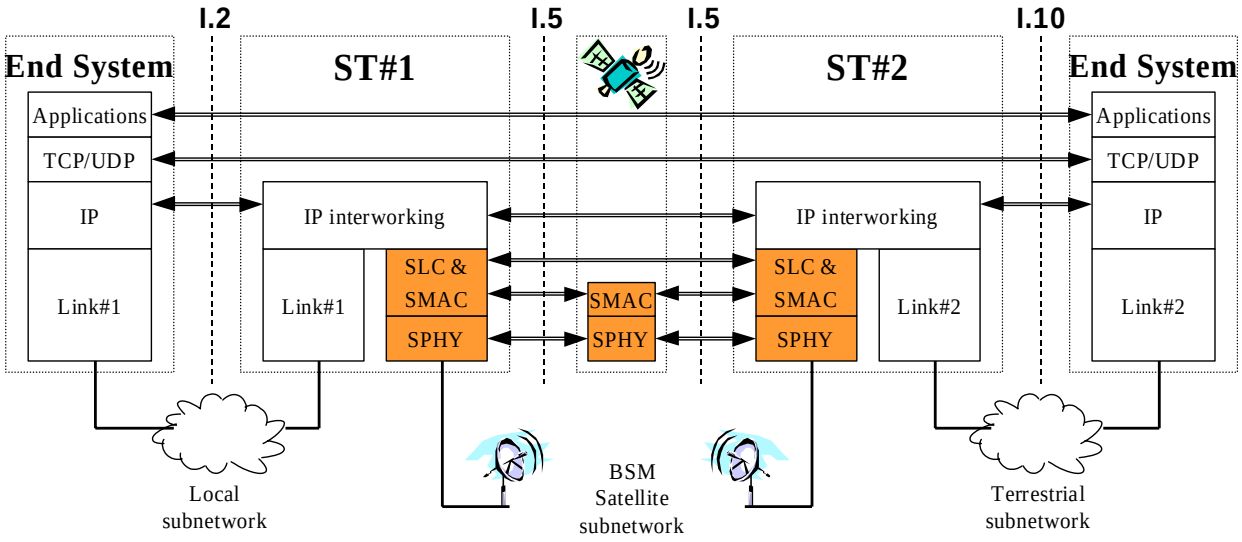

Figure 5.3.3.3: IP interworking functions

5.3.3.4 Higher layer interworking functions

A higher layer interworking function is defined as operating above the IP level of the protocol stack; i.e. it operates at a higher level than bridges or routers. A higher layer interworking function usually supports address mapping from one subnetwork to another and may also transform the data using application level interworking. A higher layer interworking function implicitly uses IP interworking; in other words IP packets are always terminated by the combined higher layer and IP interworking functions.

In the case of terrestrial subnetwork, the interworking function are typically provided by an IP router combined with additional higher level functions, where the router is used to "divert" selected IP packets up the stack into these higher level functions. Many different types of higher layer interworking function are possible, with different levels of functionality up to and including the application layer.

In the case of a satellite subnetwork, the higher layer interworking functions are broadly similar to the terrestrial gateway functions. However, satellite subnetworks introduce new satellite specific functions, such as Protocol Enhancing Proxies (PEPs). PEPs are discussed in more detail in clause 12 of the present document.

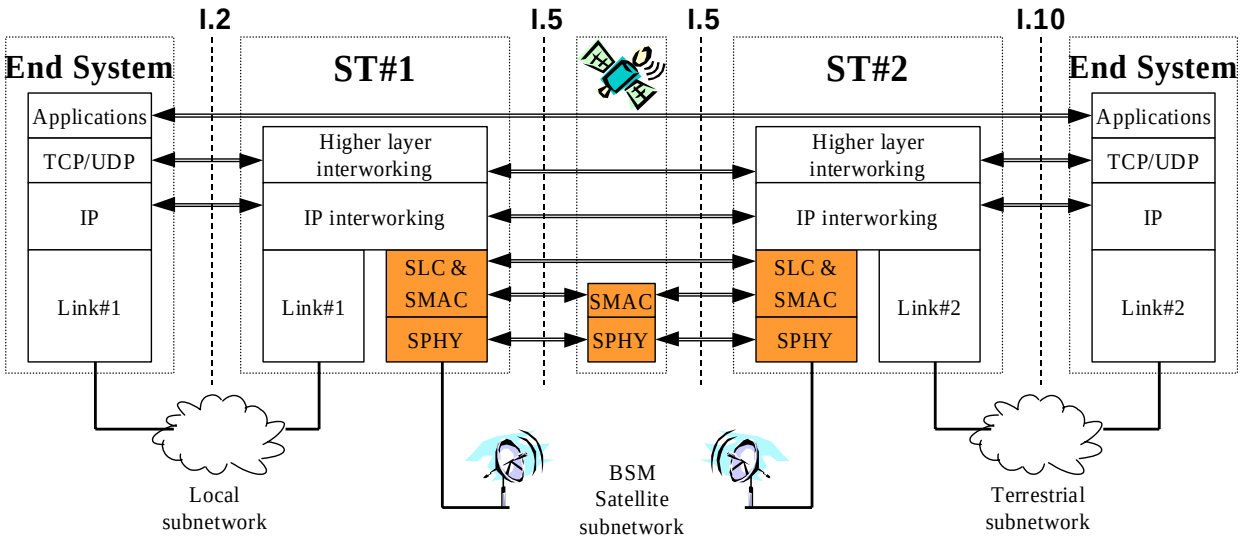

Figure 5.3.3.4: Higher layer interworking functions

6 Bearer services for transport of IP structured signals

6.1 General

IP structured signals refers to the family of network layer and higher layer protocols defined by the IETF.

The BSM network is expected to transport all IP structured signals transparently preserving the addressing and all other properties of the IP signals.

NOTE: This clause is only concerned with the basic transport properties of the BSM network. Other aspects of interworking, such as Addressing and Routing, Multicast and PEPs are discussed in separate clauses.

6.2 Network layer protocols

6.2.1 General requirements

6.2.1.1 IPv4 and IPv6

The BSM bearer services should be capable of transporting both of the following network layer protocols:

- Internet Protocol v4 (IPv4) and associated protocols [9];
- Internet Protocol v6 (IPv6) and associated protocols [10].

NOTE: IPv4 and IPv6 are the network layer protocols of the Internet TCP/IP protocol suite defined by the Internet Engineering Task Force (IETF).

6.2.1.2 IP multicast services

The BSM bearer services should be capable of transporting IP Multicast services as defined in IETF RFC 1112 [7] and IETF RFC 2236 [8].

Multicast and broadcast aspects are discussed in more detail in clause 10.

6.2.1.3 IP network services

A BSM network should efficiently support additional IP network services. Possible network services include:

- DNS services (servers, registration, etc.)
- Key management services
- Web proxies/Portals/content providers/email
- Customer management, accounting
- Helpline services
- Network management, monitoring, tuning, fault management, routing

NOTE: These network services enable the service provider to supplement and differentiate the overall BSM service for the customer.

6.2.2 Specific requirements

6.2.2.1 Maximum Transmission Units (MTUs) and IP fragmentation

IP packets have variable sizes ranging from 20 bytes to 65 535 bytes. A given subnetwork will only support packets up to a certain size known as its Maximum Transmission Unit (MTU). In order to support a variety of subnetworks, IP provides a mechanism to fragment packets that are too large for a given subnetwork. The fragments are reassembled at the destination host.

NOTE: Refer to annex A for a definition of segmentation and fragmentation.

Whereas in IPv4 fragmentation can occur at either the sending host or in an intermediate router, in IPv6 can only occur at the sending host. Because of this, IPv6 specifies a minimum MTU of 1 280 bytes. Any subnetwork with an internal packet payload smaller than 1 280 bytes must implement an internal segmentation/reassembly mechanism.

An MTU size of 1 500 bytes is currently prevalent in the Internet due to widescale deployment of Ethernet (which imposes a limit of 1 500 bytes).

6.2.2.2 BSM frame size and segmentation

The BSM subnetwork can transparently segment IP packets into internal frames [37]. The choice of the internal frame size is a difficult one, because it is necessary to choose and optimum balance between low overhead ratio due to the header and the high amount of data that will be lost if a packet is discarded. In addition to the above considerations, the following factors should also be taken into account:

- traffic type/profile; what is the IP packet length distribution?
- link speed (the time taken to transmit an MTU-sized packet over a slow link may be too large and detrimental for interactive services);
- use of tunnelling schemes, including both IPSEC tunnels and Protocol Enhancing Proxy (PEP) tunnels.

NOTE: Refer to annex A for a definition of segmentation and fragmentation.

6.2.2.3 MTU path discovery

BSM networks should support MTU path discovery [25], [37].

6.2.2.4 Maximum Segment Lifetime (MSL)

When transporting IPv4 or IPv6 packets, the BSM subnetwork should not keep and retransmit packets which have been delayed more than the IP Maximum Segment Lifetime (MSL).

NOTE: In practice this is a long time.

6.2.2.5 Reordering of packets

The BSM subnetwork should not re-order packets associated with a specific end-to-end flow. It is not necessary to provide strict in-order delivery of packets for a given flow. However, gratuitous or excessive reordering detrimentally impacts current TCP implementations.

There is no ordering requirement between packets from different hosts, or between different flows from the same pair of hosts.

6.2.2.6 Error detection

When transporting IPv4 or IPv6 packets, the BSM subnetwork should provide error detection at least as strong as the 32-bit CRC specified for HDLC [4]. The BSM subnetwork error control mechanisms should ensure that there is a low probability of undetected errors in IP packets that are delivered via the SI-SAP at the destination ST (i.e. the BSM subnetwork should only forward valid IP packets to the IP layer).

While this will achieve a very low undetected error rate in the IP packets due to transmission errors, it will not (and need not) achieve a very low packet loss rate as the Internet protocols are better suited to dealing with lost packets than with corrupted packets.

NOTE 1: This requirement is taken from the IETF PILC WG [37], [33].

NOTE 2: The IP layer may also provide filtering of packets based on the IP addresses.

6.3 Higher layer protocols

6.3.1 Transport layer protocols

6.3.1.1 User Datagram Protocol (UDP)

User Datagram Protocol (UDP) is a very simple service, that passes individual messages ("datagrams") to the IP layer for transmission. UDP is unreliable: it provides no acknowledgement of delivery and does not attempt error recovery, which (if necessary) must be undertaken by the application.

6.3.1.2 Transmission Control Protocol (TCP)

Transmission Control Protocol (TCP) provides a reliable connection oriented end-to-end transport service between hosts. TCP is designed to deliver data reliably, without errors and in sequence. TCP also contains flow control mechanisms that adjust its own behaviour in response to network conditions (notably congestion) in order to optimize overall performance.

6.3.1.3 New transport protocols

Some examples of new transport protocols are Datagram Congestion Control Protocol (DCCP) and Stream Control Transmission Protocol (SCTP).

Datagram Congestion Control Protocol (DCCP) is an unreliable transport layer that offers negotiated forms of congestion control. Maintaining many of the features of TCP, this protocol addresses many of the issues, particularly regarding the implications of satellite, and allows different congestion control to be negotiated across different parts of the link.

Stream Control Transmission Protocol (SCTP) [23], is designed to transport PSTN signalling messages over IP networks, but is capable of broader applications. SCTP is a reliable transport protocol operating on top of a connectionless packet network such as IP. The design of SCTP includes appropriate congestion avoidance behaviour and resistance to flooding and masquerade attacks.

SCTP, like TCP, provides a reliable transport service, ensuring that data is transported across the network without error and in sequence. Like TCP, SCTP is a connection-oriented mechanism, meaning that a relationship is created between the endpoints of an SCTP session prior to data being transmitted, and this relationship is maintained until all data transmission has been successfully completed.

SCTP differs by providing a number of functions that are considered critical for signalling transport, and which at the same time can provide transport benefits to other applications requiring additional performance and reliability. The core features of SCTP are multi-streaming and multi-homing.

6.3.2 Application layer protocols

6.3.2.1 Hypertext Transfer Protocol (HTTP)

Hypertext Transfer Protocol (HTTP) [24] is an application-level protocol for distributed, collaborative, hypermedia information systems. It is a generic, stateless, protocol which can be used for many tasks beyond its use for hypertext, such as name servers and distributed object management systems, through extension of its request methods, error codes and headers. A feature of HTTP is the typing and negotiation of data representation, allowing systems to be built independently of the data being transferred.

7 Performance and availability

7.1 General

Network performance and availability contribute towards the Quality of Service (QoS) as experienced by the end user. The present document only considers the network performance and availability of the BSM subnetwork (i.e. the characteristics and quality of the BSM bearer services) and the overall end-to-end performance and availability experienced by the end users will also be influenced by the other networks that are involved in the communications. In general the end-to-end performance will reflect the combined network performance of several other autonomous networks in addition to the BSM subnetwork.

7.2 Performance parameters

7.2.1 Throughput

Throughput is the parameter that defines the effective data transfer rate in bits per second (bps) for a particular service user as measured at the egress point (the exit port of the BSM network). Sharing of network capacity by a number of users reduces the throughput per user; as does any overheads added to the data by the BSM network.

From a user perspective, throughput is the primary performance parameter and different guarantees of throughput may be offered by the service provider depending on the service contract.

7.2.2 Delay

Delay is a parameter that measures the delay between the entry of a packet into one port of the satellite subnetwork and the exit of that same packet from another port of the subnetwork.

Delay manifests itself in a number of ways, including the time taken to establish a particular service from the initial user request and the time to receive specific information once the service is established. Delay has a very direct impact on user satisfaction depending on the application, and includes delays in the terminal, network, and any servers. Note that from a user point of view delay also takes into account the effect of other network parameters such as throughput.

7.2.3 Delay variation

Delay variation is a parameter that measures the differences in delay between successive packet arrivals at the exit port of the satellite subnetwork.

Delay variation is generally included as a performance parameter since it is very important at the transport layer in packetized data systems due to the inherent variability in arrival times of individual packets. However, services that are highly intolerant of delay variation will usually take steps to remove (or at least significantly reduce) the delay variation by means of buffering, effectively eliminating delay variation as perceived at the user level (although at the expense of adding additional fixed delay).

7.2.4 Transmission errors

Transmission error rates or error ratios are parameters that measure the loss or corruption of data caused by the transmission of the packet over the BSM satellite subnetwork.

Transmission errors will normally be detected through the use of a CRC (see clause 6.2.2.6) and any residual transmission errors (i.e. errors that cannot be corrected) should normally result in packet discard (packet loss).

In this context, transmissions errors are limited to the residual effects of bit errors or packet loss during transmission (i.e. only those errors that remain after applying error control techniques) and it **excludes** the effects of any degradation introduced by media coding for more efficient transmission (e.g. the use of low bit rate speech codecs for voice).

7.2.5 Availability

Availability is a parameter that measures the probability that the BSM subnetwork will provide a satisfactory service on demand. Typically, an availability statement is specified for a particular set of quality of service (QoS) parameters and these parameters can be used to define a quantitative threshold for satisfactory operation.

Availability also has a direct effect on the quality of service as perceived by the end user. Availability is often directly linked to reliability: loss of service due to propagation effects (e.g. rain fade) is equivalent to loss of service due to equipment failure from the end user perspective.

7.3 Performance objectives

7.3.1 ITU-T

ITU-T Recommendation Y.1541 [42] specifies IP performance values to be achieved internationally for each of the performance parameters defined in ITU-T Recommendation Y.1540 [43]. ITU-T Recommendation Y.1541 [42] defines six different network Quality of Service (QoS) classes and some of the performance values depend on which network (QoS) class that is agreed between the end-users and the network providers.

However, ITU-T Recommendation Y.1541 [42] applies to international end-to-end IP network paths. The network QoS classes that are defined are intended to be the basis of agreements between end-users and network service providers, and between service providers.

Note also the following points which are quoted from the scope of the Recommendation:

> "The limited number of QoS classes defined here support a wide range of applications, including the following: real time telephony, multimedia conferencing, and interactive data transfer. While the performance needs of these applications are more demanding than most, there may be other applications that require new or revised classes. Any desire for new classes must be balanced with the requirement of feasible implementation, and the number of classes must be small for implementations to scale in global networks."
>
> "The QoS objectives are applicable when access link speeds are at the T1 or E1 rate and higher."

7.3.2 IETF

In the Internet and intranets of today, in particular for multimedia applications, bandwidth and delay are important subjects. Whereas traditional Internet applications, such as HTTP, FTP or TELNET, cannot tolerate packet loss but are less sensitive to variable delays, most real-time applications show just the opposite behaviour, meaning they can compensate for a reasonable amount of packet loss but are usually more critical towards high and/or variable delays. This means that without any bandwidth control, the quality of these real-time streams depends on the bandwidth that is instantaneously available.

The use of the Internet Protocol (IP) in broadband satellite multimedia (BSM) networks presents challenges not usually faced with in terrestrial wireline networks. Bandwidth is scarce, hence has to be managed carefully, delay is usually high and can be made worse by the use of bandwidth on demand, and availability of network resources can be low due to weather events. As a consequence, BSM networks and especially the BSM quality of service and performance protocols will have to alleviate delay, limited bandwidth, and varying capacity and build on the strength of BSM in terms of coverage and natural multicasting to ensure that BSM subnetworks become an integral part of the Internet.

IP over satellite goes back to the old Arpanet and Satnet experiments (1970s and 1980s), while those were essentially trunking and point to point they nevertheless demonstrated that satellites did not break IP. In the mid 1990's the IETF group TCP over Satellite (TCPsat) investigated how TCP could be used in BSMs without special changes to the Internet Protocol (IETF RFC 2488 [32]). Since then, the Performance Implications of Link Characteristics (pilc) working group [37] has studied how IP performance can be improved using proxies (IETF RFC 3135 [20]) and how link designers and network engineers need to interact to ensure maximum performance of IP over heterogeneous networks [33]. This work is continuing in the IETF pilc working group and the bulk of this work is applicable to BSM systems.

7.4 Characteristics of satellite links

7.4.1 TCP delays

TCP operation over satellite links has been extensively studied in IETF. IETF favour end-system solutions, such as modifications to the TCP stacks.

An ST contains several different elements which can introduce delays and thereby affect the operation of end-to-end TCP:

- Protocol Enhancing Proxies (PEPs) as discussed in clause 12.
- The air interface Bandwidth on Demand schemes.
- Various queues and buffers.

7.4.2 Bit error ratio

The bit error ratio (BER) is defined as the ratio of the number of transmission errors (erroneous or lost bits) to the total transmitted bits.

Typically the transmissions over a satellite link are protected using a combination of error correction and detection and the resulting protected link is error free during normal operation, lapsing into errors only during occasional outages. The objective is to achieve the "error free" transmission state for as long as possible, taking into account that the rest of quality parameters (user bandwidth, delay, etc.) are the responsibility of the higher layers, the application provider or the terrestrial network operator.

The atmospherics conditions and the corresponding attenuation are the primary cause for the outages of the link. The procedure to calculate the rain and atmospheric attenuation is reflected in ITU-R Recommendations P.618 [38], P.837 [39], P.838 [40] and P.839 [41].

8 Quality of Service (QoS)

8.1 Overview of QoS

8.1.1 QoS definition

Quality of Service (QoS) is the collective effect of service performance which impacts the degree of satisfaction of a user of the service.

QoS is to the ability of a network element (e.g. an application, host or router) to have some level of assurance that its traffic and service requirements can be satisfied. To enable QoS requires the co-operation of all protocol layers from top-to-bottom, as well as every network element from end-to-end. Any QoS assurances are only as good as the weakest link in the "chain" between source and destination.

8.1.2 QoS architecture

This clause discusses Quality of Service (QoS) as applied to the BSM bearer services.

Network Services are considered end-to-end, this means from an End System (ES) to another ES. An end-to-end Service may have a certain Quality of Service (QoS) which is provided for the user of a network service. It is the user that decides whether he is satisfied with the provided QoS or not.

To realize a certain network QoS a bearer service with clearly defined characteristics and functionality is to be set up from the source to the destination of a service.

A bearer service includes all aspects to enable the provision of a contracted QoS. These aspects are among others the control signalling, user plane transport and QoS management functionality. A BSM bearer service architecture is depicted in figure 8.1.2 (reproduced from the BSM Service and Architecture report [3]. Each bearer service on a specific layer offers its individual services using services provided by the layers below.

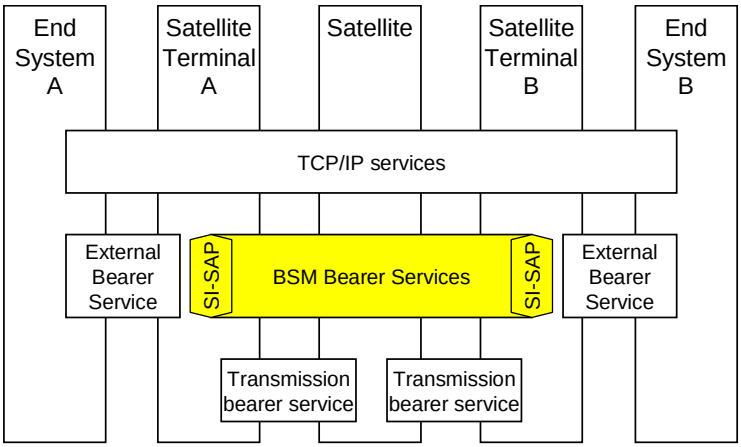

Figure 8.1.2: BSM bearer services

8.1.3 End-to-End QoS

On its way from one End System (ES) another ES the traffic has to pass different bearer services of a series of network(s) and the End-to-End-Service used by the ES is determined by the combination of the BSM bearer services and those External Bearer Services. As the End-to-End Service is conveyed over several networks (not only BSM) it is outside the scope of the present document.

The External Bearer Services are not further elaborated in the present document: these bearer services may be using several network services, e.g. an ES - ST Local Bearer Service, or an External Bearer Service provided by terrestrial transit networks or even another BSM Bearer Service.

8.1.4 BSM Quality of Service

BSM Quality of Service (BSM QoS) is the QoS that applies to the BSM Bearer Services. These are the services that the BSM network operator offers and it is these BSM bearer services that provide the BSM QoS.

QoS does not create bandwidth. It is not possible for the network to give what it does not have, so bandwidth availability is a starting point. QoS only manages bandwidth according to application demands and network management settings, and in that regard it cannot provide certainty if it involves sharing. Hence, QoS with a guaranteed service level requires resource allocation to individual data streams. A priority for QoS designers has been to ensure that best-effort traffic is not starved after reservations are made. QoS-enabled (high-priority) applications must not disable the mundane (low-priority) Internet applications.

When defining the BSM QoS, the restrictions and limitations of the air interface have to be taken into account. It is not reasonable to define complex mechanisms as have been in fixed networks due to different error characteristics of the air interface. The QoS mechanisms provided in the BSM network have to be robust and capable of providing reasonable QoS resolution.

8.2 IP Quality of Service

8.2.1 General

There is no single IP QoS service model. Broadly speaking two approaches exist: Integrated Services (Intserv) [9] and Differentiated Services (Diffserv) [10]. Combinations of the two have also been proposed. IETF RFC 2990 [11] represents the present understanding of the challenges in providing a QoS architecture for the Internet.

Most current IP networks do not support either of these approaches and instead adopt a "Best Effort" service.

8.2.2 Best-effort

The basic IP protocol stack provides only one QoS, which is called best-effort. Packets are transmitted from point to point without any guarantee for a special bandwidth or minimum time delay. This means that all requests have the same priority and there is no possibility to make bandwidth reservations for specific connections or to raise the priority for special requests. Therefore, new strategies were developed to provide predictable services for the Internet. Today, there are two main approaches to bring QoS to the Internet and IP based internetworks: Integrated Services and Differentiated Services.

8.2.3 IP Integrated Services (Intserv)

Integrated Services (Intserv) bring enhancements to the IP network model to support real-time transmissions and guaranteed bandwidth for specific flows. In this case, we define a flow as a distinguishable stream of related datagrams from a unique sender to a unique receiver that results from a single user activity and requires the same QoS. For example, a flow might consist of one video stream between a given host pair. To establish the video connection in both directions, two flows are necessary. Each application that initiates data flows can specify which QoS are required for this flow. If the video conferencing tool needs a minimum bandwidth of 128 kbps and a minimum packet delay of 100 ms to assure a continuous video display, such a QoS can be reserved for this connection.

IP Integrated Services (Intserv) (IETF RFC 1633 [9]) provide fine-grained service guarantees to individual flows. Flows are identified by a flow specification (flowspec), which creates a stateful association between individual packets by matching fields in the packet header. Bandwidth is reserved for the flow, and appropriate traffic conditioning and scheduling is installed in routers along the path.

The ReSerVation Protocol (RSVP) (IETF RFC 2205 [12]) (IETF RFC 2210 [13]) is usually, but not necessarily, used to install flows. Intserv defines two services, in addition to the Default (best effort) service:

- Guaranteed Service (GS) (IETF RFC 2212 [15]) offers hard upper bounds on delay to flows that conform to a traffic specification (TSpec). It uses a fluid flow model to relate the TSpec and reserved bandwidth (RSpec) to variable delay. Non-conforming packets are forwarded on a best-effort basis.

- Controlled Load Service (CLS) (IETF RFC 2211 [14]) offers delay and packet loss equivalent to that of an unloaded network to flows that conform to a TSpec, but no hard bounds. Non-conforming packets are forwarded on a best-effort basis.

Intserv requires installation of state information in every participating router, and if this is not present in every router along the path, performance guarantees cannot be made. This, along with RSVP processing and the need for usage-based accounting is believed to have scalability problems, particularly in the core of the Internet (IETF RFC 2208 [19]).

8.2.4 IP Differentiated Services (Diffserv)

Differentiated Services Differentiated Services mechanisms do not use per-flow signalling, and as a result, do not consume per-flow state within the routing infrastructure. Different service levels can be allocated to different groups of users, which means that all traffic is distributed into groups or classes with different QoS parameters. This reduces the maintenance overhead in comparison to Integrated Services.

IP Differentiated Services (Diffserv) (IETF RFC 2475 [10]) provides a "toolkit" offering coarse-grained controls to aggregates of flows. Diffserv in itself does not provide QoS guarantees, but can be used to construct services with QoS guarantees across a Diffserv domain. It attempts to address the scaling issues associated with Intserv by requiring state awareness only at the edge of a Diffserv domain. At the edge, packets are classified into flows, and the flows are conditioned (marked, policed or shaped) to a Traffic Conditioning Specification (TCS).

A Diffserv Codepoint (DSCP), identifying a per-hop behaviour (PHB) is set in each packet header. The DSCP is carried in the DS-field, subsuming six bits of the former TOS byte of the IP header (IETF RFC 2474 [16]). The PHB denotes the forwarding behaviour to be applied to the packet in each node in the Diffserv domain. Although there is a "recommended" DSCP associated with each PHB, the mappings from DSCPs to PHBs are defined by the DS-domain. In fact, there can be several DSCPs associated with the same PHB.

Diffserv presently defines three PHBs. The class selector PHB (IETF RFC 2474 [16]) replaces the IP precedence field of the former TOS byte. It offers two relative forwarding priorities:

- The Expedited Forwarding (EF) PHB (IETF RFC 2598 [18]) guarantees that packets will have a well-defined minimum departure rate which, if not exceeded, ensures that the associated queues are short or empty. EF is intended to support services that offer tightly bounded loss, delay and delay jitter.

- The Assured Forwarding (AF) PHB group (IETF RFC 2597 [17]) offers different levels of forwarding assurances for packets belonging to an aggregated flow. Each AF group is independently allocated forwarding resources. Packets are marked with one of three drop precedences, such that those with the highest drop precedence are dropped with lower probability than those marked with the lowest drop precedence. DSCPs are recommended for four independent AF groups, although a DS domain can have more or fewer AF groups.

8.3 IP transfer capabilities

8.3.1 General

An IP transfer capability is a set of network capabilities provided by IP based networks to transfer IP packets. For each IP transfer capability, the service model, traffic descriptor, conformance definition and any QOS commitments are defined. An IP transfer capability is supported by a set of traffic control and congestion control functions.

In order to offer multiple classes of QOS to multiple applications and to optimize the usage of network resources, IP based networks should be capable of providing multiple transfer capabilities.

Three IP transfer capabilities are defined in ITU-T Draft Recommendation Y.iptc [44]:

- Dedicated Bandwidth (DBW) IP transfer capability

- Statistical Bandwidth (SBW) IP transfer capability

- Best Effort (BE) IP transfer capability

This set of IP transfer capabilities are based on current IP service models and this set may be extended in the future.

8.3.2 Dedicated Bandwidth (DBW) transfer capability

The Dedicated Bandwidth (DBW) transfer capability is intended to support applications with stringent delay requirements. It aims to support the guaranteed and timely delivery of IP packets along the end-to-end path of the network.

The DBW transfer capability strives for compatibility with the Guaranteed Service; IETF RFC 2212 [15] and the end-to-end services based on the Expedited Forwarding per-hop behaviour; IETF RFC 2598 [18].

The commitment made by the network is that the negotiated IP QOS is assured to all IP packets when all packets are conforming to the conformance tests. The DBW user should expect that (possibly all) non-conforming packets be discarded by the network.

8.3.3 Statistical Bandwidth (SBW) transfer capability

The Statistical Bandwidth (SBW) transfer capability is intended to support applications, which do not have stringent delay requirements. It aims to support the guaranteed delivery of IP packets along the end-to-end path of the network.

The SBW transfer capability strives for compatibility with the Controlled Load Network Element Service; IETF RFC 2211 [14] and the end-to-end services based on the Assured Forwarding per-hop behaviour; IETF RFC 2597 [17].

The SBW transfer capability provides a specified sustainable rate (Rs) for non-real time applications with limited burst duration with the expectation that excess traffic will be delivered within the limits of available resources. The SBW capability may be associated with a specified packet loss commitment.

8.3.4 Best effort (BE) transfer capability

The best effort IP transfer capability is intended to support applications which do not have stringent loss or delay requirements.

The service model for the best effort (BE) IPTC requires that available resources be used for forwarding packets of best effort flows. Even though there are no QOS commitments specified, the expectation is that packets be delivered provided that sufficient resources are available.

8.4 BSM QoS to IP QoS interworking

8.4.1 Background

The Internet Protocol (IP), and the architecture of the Internet itself, is based on the simple concept that datagrams with source and destination addresses can traverse a network of (IP) routers independently, without the help of their sender or receiver. The Internet was historically built on the concept of a dumb network, with smarts at either end (at the sender and receiver).

There is a price to pay for this simplicity, however. The reason IP is simple is because it does not provide many services. IP provides addressing, and that enables the independence of each datagram. IP can fragment datagrams (in routers) and reassemble them (at the receiver), and that allows traversal of different network media. But IP does not provide reliable data delivery. Routers are allowed to discard IP datagrams en route, without notice to sender or receiver. IP relies on upper-level transports (e.g. TCP) to keep track of datagrams, and retransmit as necessary. And these "reliability" mechanisms can only assure data delivery; neither IP nor its high-level protocols can ensure timely delivery or provide any guarantees about data throughput. IP provides what is called a "best effort" service. It can make no guarantees about when data will arrive, or how much it can deliver.

This limitation has not been a problem for traditional Internet applications like web, email, file transfer, and the like. But the new breed of applications, including audio and video streaming, demand high data throughput capacity (bandwidth) and have low-latency requirements when used in two-way communications (i.e. conferencing and telephony). Public and private IP Networks are also being used increasingly for delivery of mission-critical information that cannot tolerate unpredictable losses. Unlike "pure virtual circuit" technologies like ATM and Frame Relay, IP does not make hard allocations of resources. This provides much more efficient use of the available bandwidth, and it is more flexible also. Typical network traffic is bursty rather than continuous. IP is datagram-based so it uses the available bandwidth most efficiently, by sharing what is available as needed. This also allows IP to adapt more flexibly to applications with varying needs. However, it also leads to some unpredictability in service. The capacity to tolerate this unpredictability relates to the level of guarantee they require.

8.4.2 Mapping IP QoS to BSM QoS

IP-based applications do not directly use the BSM bearer services but they use IP QoS definitions and attributes, which are mapped to BSM QoS attributes at the SI-SD interface. In the case of interworking between IP networks and a BSM network for the transport of IP-based applications, the selection of the BSM QoS and associated QoS attribute values is derived from the Internet QoS attributes.

The challenge of QoS and performance in BSM is to define how the interface between the BSM and the other subnetworks of the Internet is going to be managed.

Currently there are two main Internet QoS concepts, namely Integrated Services and Differentiated Services and the BSM QoS should be capable of supporting both of these concepts in addition to the basic Best_Efforts IP service.

IP based QoS models should be supported for both Integrated Services (IntServ) signalled by RSVP (IETF RFC 2205 [12]) and Differentiated Services (6-bit QoS attribute on each IP packet, DiffServ). These IP based QoS shall be mapped to BSM QoS by a network element at the border of the BSM network, such as the Satellite Access Function (SAF) or Satellite Gateway Function (SGF) as defined in the BSM services and architectures report [3].

RSVP support would require flow establishment, and possibly aggregation of flows, within the BSM network. Differentiated services would require that there is either one QoS profile for each traffic type or alternatively the priority and traffic type information is included in the data packets.

The BSM QoS model should also permit the QoS attributes to be used to provide different grades of service to different users. In other words, the QoS attribute values for particular flows could be determined wholly or partly as part of the subscription. For example, business users may wish to subscribe to a higher grade of service than residential users, and be prepared to pay a higher fee in return for higher performance.

8.4.3 Functional model for BSM QoS

A functional model for a possible implementation of BSM QoS is illustrated in figure 8.4.3. This contains two main functional components:

- C-plane functions that establish BSM bearer services in response to user demands. This includes BSM bearer service control above the SI-SAP and the related bearer service manager below the SI-SAP.
- U-plane functions that operate on the individual packets. This includes Packet classification and packet conditioning functions above the SI-SAP and Admission control functions below the SI-SAP.

A subscription control function is also included in figure 8.4.3. This function provides the authorization for all user transactions; both C-plane and U-plane.

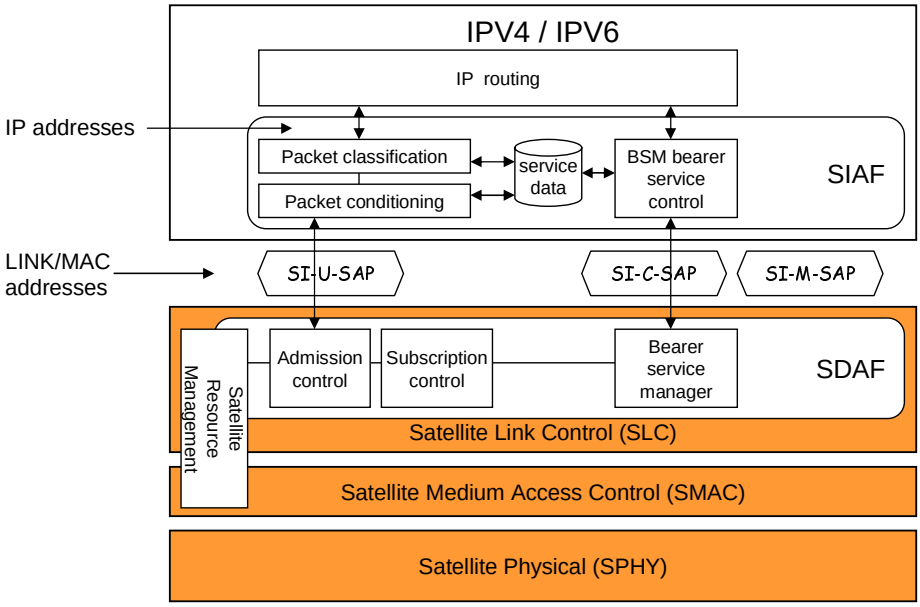

Figure 8.4.3: QoS functional model

9 Routing and Addressing

9.1 General

In the IP world, when a packet comes into a router the following action takes place:

- The router examines the destination IP address to determine if the router is the destination.
- If so, the router passes the IP packet "up the stack" to the appropriate application based on the protocol field.
- If not, the router decrements the TTL (or Hop Count) field (and discards the IP packet if zero).
- The router examines the destination IP address to determine if the destination host belongs to the subnets associated with any of the router's interfaces.
- If so, the router performs address resolution of the destination host IP address as described below.
- If not, the router determines the "next hop" IP address from the route table.
- Address resolution of the next hop (or destination host) IP address is performed. Address resolution determines the MAC layer address associated with the next hop IP address (an example is the Ethernet MAC address).
- The IP datagram is passed "down the stack" where it is encapsulated with Link Layer and MAC layer headers and the resulting Packet is then submitted to the router interface as determined by the route table.

The BSM network should be designed such that this normal IP router operation can operate over the BSM satellite subnetwork as though it were a terrestrial network, as far as possible. In particular:

- Address resolution should be capable of determining the Satellite MAC layer address associated with the next hop IP address, where appropriate.
- The IP datagram should be passed "down the stack" to the satellite interface where it is encapsulated with Satellite Link Layer and Satellite MAC headers.

9.2 Address resolution

This clause discusses how the BSM system determines the IP address and MAC address of the "next hop" in order to forward IP datagrams to a router at another ST.

9.2.1 General

Address Resolution (AR) is the means by which a network layer (IPv4 or IPv6) address is resolved to a link layer (satellite MAC or Ethernet) address. Address resolution is performed after the router interface is determined and makes use of an AR cache, which keeps AR entries for resolving the network address. The AR cache may be populated by a network protocol associated with the router interface; this clause will identify the protocols used for AR. In addition, static AR entries may also be configured under certain conditions.

The satellite interface may make use of AR broadcasts to update the AR cache at each ST, thereby improving the efficiency when compared with point-to-point routing protocols.

9.2.2 Reference model

Figure 9.2.2 shows a reference model for the Addressing and Routing functions at the Satellite Independent interface (SI_SAP). The model defines two components:

- The address resolution function in the C-plane. This function is used to determine the satellite link address when the address translation is unknown. The results of address translation are stored in the cache for future use.
- The satellite address mapping function in the U-plane. This function maps the IP address to the corresponding satellite link address (e.g. a Satellite MAC address). This function makes use of an address cache which stores the address pairs.

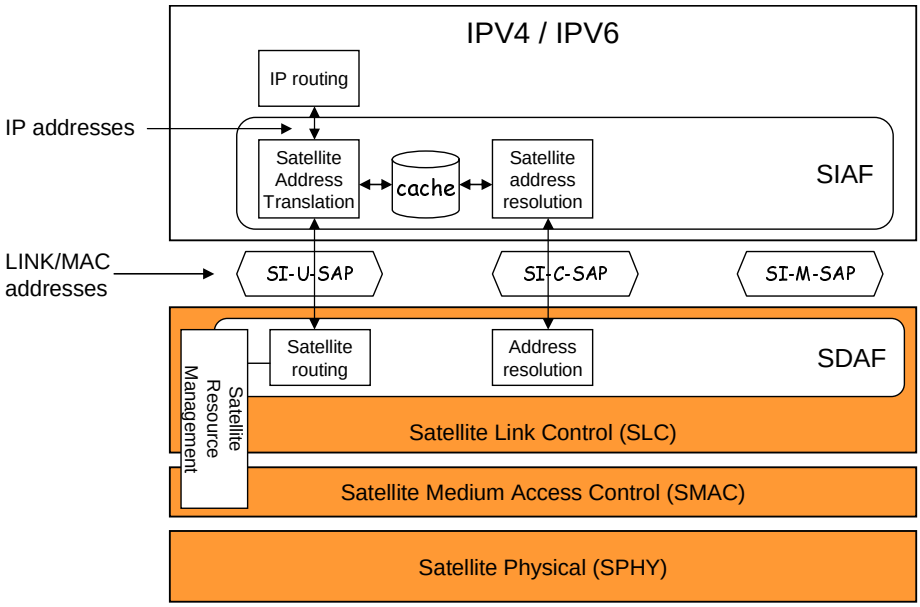

Figure 9.2.2: Address and Routing reference model

An IP (IPv4 or IPv6) network address is assigned to the upper layers above the SI-SAP, while a MAC address is assigned to the lower layers below the SI-SAP. Resolution between these two addresses is performed when an ST intends to send IP packets across the satellite link to a different ST as the next hop IP router.

Note that an IP address is assigned to the satellite air interface and an additional IP network addresses (IPv4 or IPv6) will normally be required for the other (wired) IP interfaces (e.g. Ethernet ports).

9.2.3 External interfaces

For the external (wired) interfaces, the ST shall use standard Address Resolution protocols.

For example, for an Ethernet interface, ARP (Address Resolution Protocol), IETF RFC 826 [21] shall be used for resolving IPv4 addresses and ND (Neighbour Discovery), IETF RFC 2461 [22], for resolving IPv6 addresses.

9.2.4 SI-SAP interface

Satellite specific Address Resolution protocols are assumed internal to the satellite subnetwork. AR protocols that are used for external network are generally not suitable since these protocols may generate excessive signalling traffic. Moreover, standard IP routing protocols are usually designed to allow multiple routes, differentiated only by a "cost" metric (e.g. OSPF). By contrast, the satellite routing requires constrained routing, which takes account of the special properties of the satellite paths.

However, it should be possible to support standard Address Resolution services via the SI-SAP that allow the external IP routers to access Address Resolution information without participating directly in the internal process.

Because the BSM network is a shared resource where many subnets all use a single communication link, the next-hop IP addresses assigned to BSM users are not guaranteed to be unique. However, uniqueness is required in order to resolve to the satellite MAC address. The IP-addresses should therefore be mapped into a unique addressing by adding appropriate "context" fields to the IP-address which uniquely identify the user/subnet. The combined IP-address plus context fields can then be to perform address resolution for the satellite interface.

9.2.5 Satellite specific address resolution protocols

Through the SI-SAP, address resolution services can make use of an internal satellite specific address resolution protocol (S-ARP). Such protocol can make use of the inherent broadcast nature of the satellite links without introducing the overhead of Ethernet ARP or IPv6 Neighbour Discovery (ND).

A possible solution would consist of a centralized server, co-located with the Network Control Centre (NCC), which receives unicast address resolution requests and replies instead of the target node. The link layer address of the server is given to the ST during the registration of the ST user to its Network Operator.

Instead of using a centralized server the address resolution database could be distributed among a number of servers. This solution is more flexible as it allows the presence of more than one Network Operators in a same satellite link and allows address resolution across Network Operators.

For example, in figure 9.2.5 ST1 wants to send a datagram to ST2. ST1 knows the IP address of ST2 (x) but does not know its link layer address (d). Then ST1 sends a unicast address resolution request to its address resolution server. Such server is accessed through a hub terminal. The server processes this request and sends the response including the link layer address d of ST2 to ST1. Then ST1 caches the advertised address resolution entry and sends its datagram to ST2 using ST2's link layer address.

The kind of link layer addresses or identifiers used in the BSM network at the access layer is important as it strongly impacts the overhead at layer 2 as well as the S-ARP signalling exchange load. Such identifier would be analogous to an Ethernet MAC address, an ATM VPI/VCI or a MPLS Label. But there needs to be a direct mapping between this identifier and the identification fields existing in the MAC layer's PDU formats used: for instance a VPI/VCI for a ST emitting according to DVB-RCS ATM's profile. Such identifier should also be easy to process by satellite OBP in order to switch packets.

A new labelling scheme could be developed where a same layer 2 label would be shared by several ST's at emission and at reception. Such label could be mapped to a VPI or a MPEG PID. Sharing a same label between several ST's at emission and reception allows, provided ST's can perform IP filtering, a connectionless transfer mode. This helps reducing signalling load and avoiding complex signalling functions (such as ATM UNI). It also facilitates multicast, and especially multi-source multicast.

If STs have no layer 3 functions, the same labelling and S-ARP principles should be applicable.

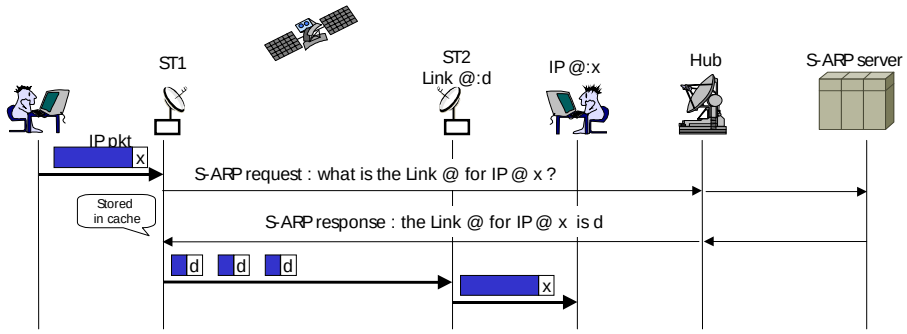

Figure 9.2.5: Satellite ARP

9.3 Routing

9.3.1 General

Routing determines the network layer (IPv4 or IPv6) address of the next hop that a network layer packet must be sent.

An ST should have the capability to perform IP routing functions. As a minimum the ST should deal with static routes to enable subnets on the satellite side of the terminal to be owned by the satellite operator or service provider but subnets on the user side to be owned by the user. Support for dynamic routes will also be needed in some cases.

9.3.2 Static routing

This clause describes the elements needed for forwarding IP packets through the BSM system using static or default routing. Static routing may be applicable to both the terrestrial and satellite interfaces.

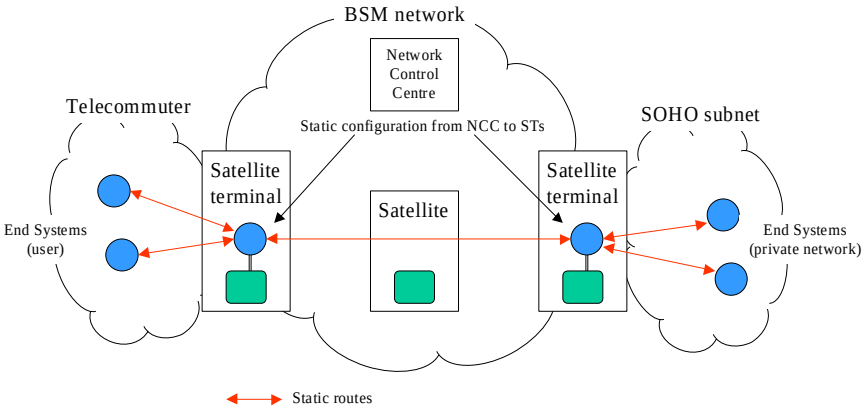

Figure 9.3.2: Static routing

Figure 9.3.2 shows a telecommuter PC (or PCs) being supported by an ST using static routing. In this case, the service provider (or equivalent entity) configures a static route into the ST supporting the telecommuter whose destination address matches the SOHO subnet and whose next hop points will contain the network address of the SOHO ST. The Network Control Centre (NCC) will provide the service provider with the ability to configure these static route entries which will contain the destination IP subnet, a subnet mask, the additional "Context" fields (see above). When packets come in the terrestrial interface of the ST supporting the telecommuter, a route table lookup is done. If the packet is destined for the SOHO network, then the static route configured by the NCC will be selected.

9.3.3 Dynamic routing

This clause describes the forwarding of IP packets through the BSM system using dynamic routing. Dynamic routing may be applicable to both the terrestrial and satellite interfaces. With dynamic routing, the forwarding table of the ST is configured and updated dynamically thanks to a routing protocol. This requires that a routing protocol is implemented in each ST, in order that STs exchange routing information on the satellite link, and fill and update their routing tables. Each ST advertises on satellite links destinations that it can access by a terrestrial way.

The routing protocol should be chosen carefully, according to several criteria:

- performance of the protocol in a satellite system context;
- Interior Gateway Protocol (IGP) such as OSPF, RIP, IS-IS or Exterior Gateway Protocol (EGP) such as BGP, depending on the fact if the STs belong to the same Autonomous System or not.

When packets enter the terrestrial interface of the ST, a route table lookup is done. If their destination address matches one entry, packets are sent to the corresponding next hop, else they are sent to the default next hop.

10 Multicast and Broadcast

10.1 General

By virtue of their wide-coverage area, BSM systems can be particularly effective when used to provide two basic categories of service: multicast and broadcast.

Here we use the term "multicast" to indicate an addressed service and broadcast to indicate an unaddressed (or "all-stations") service. Both services feature a point-to-multipoint or multipoint-to-multipoint topology.

Multicasting and broadcasting can both be used for real-time or non real-time services. Non-real time services may use retransmissions to provide a good quality of service via satellite. Local caching may be used as an addition to the non-real time services: for a cached service the traffic is received by the ST and stored in local terminal memory (e.g. a hard disc) and the user can access the service at any time after successful reception.

Multicast services may offer a return channel from the user to the network. This can be used to provide selective acknowledgements, or to provide support for IP multicast signalling (join and leave instructions) as discussed below.

10.2 Reference models

10.2.1 IP multicast model

A reference model for IP multicasting via a BSM network is illustrated in figure 10.2.1. The figure illustrates several different functions that can perform replication of the multicast packets:

- IP replication by the IP router in the source ST;
- BSM replication in the source ST node;
- BSM replication in the satellite;
- IP replication by the IP router in the destination ST.

BSM replication refers to replication within the BSM satellite subnetwork, using internal functions that are specific to the satellite subnetwork. For example, packet replication at this level may be used for spot beam systems to replicate the data into more than one spot beam.

A BSM system can use a combination of these replication functions, with the chosen method(s) being determined by the satellite specific capabilities (on-board processing, spot beam etc). The replication method(s) may also be selected dynamically depending on the service required; for example depending on the number of intended destinations.

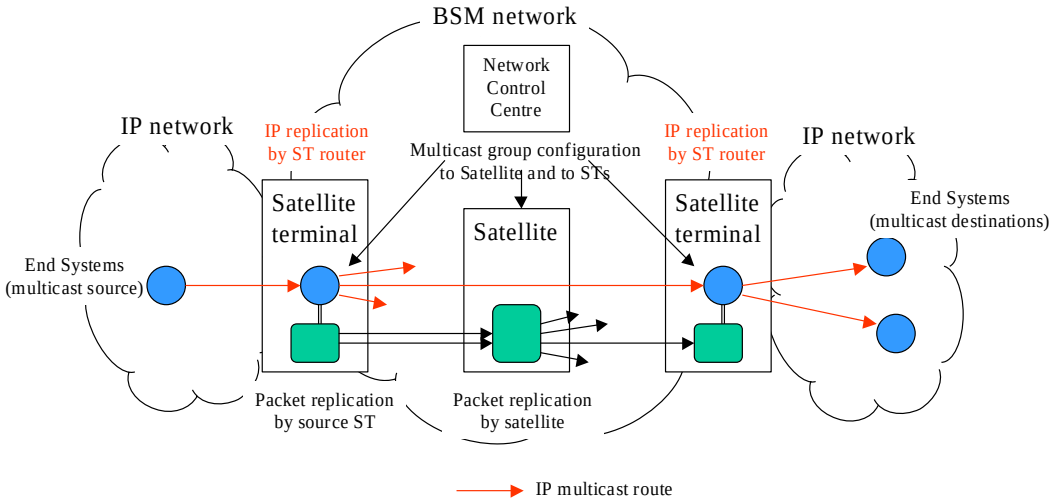

Figure 10.2.1: IP Multicast model (forward direction)

10.2.2 Addressing functional model

The functional model for addressing contains three main components as illustrated in figure 10.2.2:

- Multicast addressing resolution functions in the C-plane. These are a subset of the general address resolution functions defined in clause 9.
- Multicast group management functions in the C-plane. These are additional C-plane functions that translate the IP group management protocols into internal satellite group management protocols.
- The satellite address mapping function in the U-plane. These are a subset of the general addressing mapping functions defined in clause 9.

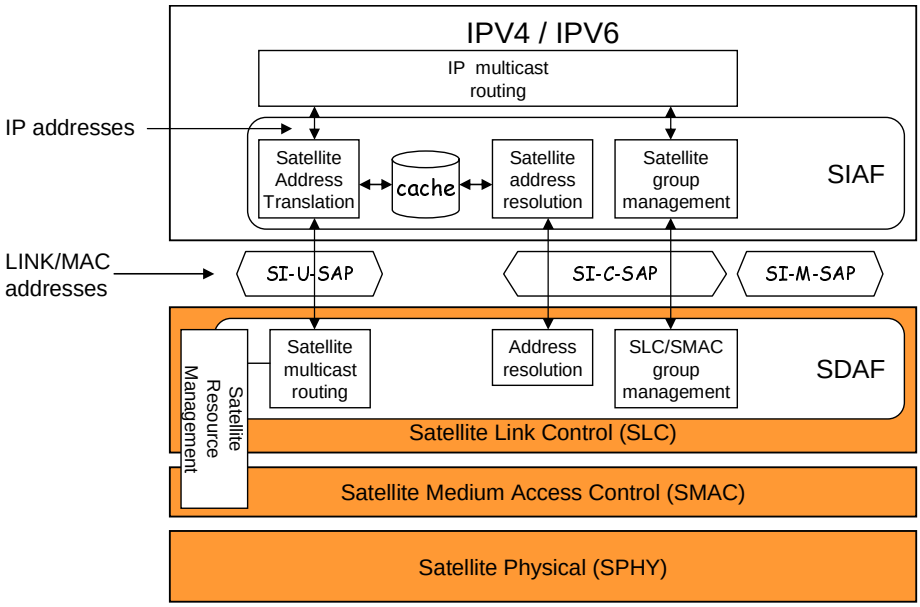

Figure 10.2.2: Addressing functional model

The addresses that are used at each layer depend on where the multicast replication functions are located.

10.2.3 Replication functional model

There are several alternative positions for the multicast replication functions as illustrated in figure 10.2.3.

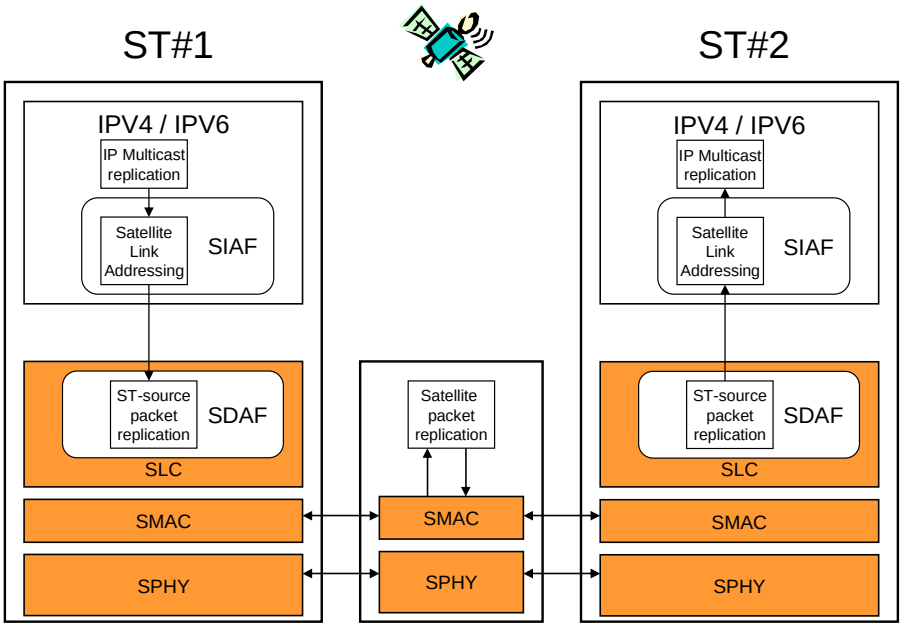

Figure 10.2.3: Replication functional model

In the general case, IP multicast addresses are used, together with IP multicast replication. A LINK/MAC address is then added and the resulting multicast packet is passed to the SLC layer. At this layer additional replication may be needed depending on the LINK/MAC address used and depending on the satellite network. The satellite network replication can be divided into two main techniques based on the transmission service that is used:

- broadcast-transmission techniques; where the multicast packets are not replicated and a transmitted once in a single broadcast beam;
- multicast-transmission techniques; where the multicast packets are replicated and transmitted in one or more beams, where each individual transmission may be addressed to multiple destination.

The multicast transmission technique may use a combination of source ST replication, satellite replication and destination ST replication to minimize the total satellite capacity that is used.

10.3 IP Multicast functions

IP Multicast Services are defined in IETF RFC 1112 [7] and IETF RFC 2236 [8]. These services differ from the generic BSM multipoint services (as defined above) in that the transmission is from one or more users to a single group address (a specific IP address). None or more destinations may join the group to request copies of the packets with a specified group address. Users (senders) operate independently of destinations. There is no implied order in which users and destinations join and hosts may join (subscribe) and leave (unsubscribe) as and when they wish. A user (sender) need not join a group; if it does not, it will not receive a copy of the packets.

In IP multicast, the network delivers a packet (or set of packets) using a class D IP destination address to identify a group. End hosts indicate which groups (or a combined source IP address and group destination address) to the network, and hosts may subscribe to groups (join) and unsubscribe (leave) as and when they wish. There is no explicit list of destination addresses; therefore this information must be derived from an IP group membership protocol (e.g. IGMP) or via an IP multicast routing protocol (e.g. PIM-SM).

These IP multicast signalling protocols have been specified for IP terrestrial networks (local network or intra-domain or inter-domains). It seems necessary to either adapt them or to suggest some special configurations for a BSM network to take account of the specific properties of a BSM network. The following issues are noted:

- A BSM network has a longer delay than a terrestrial network.

- The number of routers connected to the satellite link is greater than the number of routers on a LAN and is also greater than the number of neighbour routers on a terrestrial domain.
- The bandwidth on BSM network needs to be efficiently used by services. It is important to prevent multicast signalling protocols from flooding the satellite link with messages.

Therefore an adaptation is required on entities implementing multicast signalling protocols on a BSM network (STs themselves or entities directly connected to STs). However a BSM network should offer a standard multicast interface towards IP networks to be compliant with multicast terrestrial networks.

Any BSM network wishing to support IP multicast efficiently should support multicast natively at the link layer. In addition, a mapping function would have to be developed to translate from IP multicast addresses to link layer multicast addresses.

BSM networks may wish to offer both static and dynamic multicast groups between Satellite Terminals (on the air interface). For static multicast no multicast signalling protocol is used between Satellite Terminals whereas a signalling protocol is required for dynamic multicast groups.

10.3.1 Static multicast groups

Static multicast groups are groups that are pre-configured by management. A ST can either accept or discard data from the groups to which it is subscribed.

Static multicast groups may be permanent, or scheduled. Scheduled multicast groups are valid for a specific period: this may be a single event or a regular event.

Static multicast groups should include support for IP streaming services.

10.3.2 Dynamic multicast groups

Dynamic multicast groups are groups that allow the ST to elect to join or leave the group at any time, subject to network capability and the ST subscription.

Dynamic multicast groups should include support for IP multicast.

10.3.3 Multicast addressing

Multicast services use specific IP address in Class D to communicate. The address range defined for multicast is 224.0.0.0 - 239.255.255.255. Some multicast addresses are reserved to multicast signalling protocols on a local network such as group membership protocol (IGMP) and routing protocols, Session Announcement Protocol (SAP).

A multicast service is defined by an IP multicast address as the destination address in the packet. The source address may be an IP unicast address (e.g. the source address of the multicast server).

The multicast routing and forwarding functions are implemented at the IP Layer. However, to improve the efficiency of the forwarding function in ST, it is recommended to use filtering function at the lower layers; this is often implemented in hardware. Thus the lower layer will stop as far as possible uninteresting multicast flow and charge the forwarding function with the minimum of uninteresting flows.

A correspondence is therefore required between IP multicast address and lower layer multicast addressing. This correspondence may be systematic, announced, or configured.

With a systematic correspondence, the lower layer address is calculated in function of the IP multicast address.

EXAMPLE 1: Ethernet/FDDI Mac address Mapping as defined in IETF RFC 1112 [7].
"An IP host group address is mapped to an Ethernet multicast address by placing the low-order 23-bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01-00-5E-00-00-00 (hex). Because there are 28 significant bits in an IP host group address, more than one host group address may map to the same Ethernet multicast address." (IETF RFC 1112 [7]).

With an announced correspondence, a Hub terminal or the Network Control Centre announces by periodically sending tables within correspondences.

EXAMPLE 2: Multicast Mapping Table (MMT).
This is not a standard but corresponds to a current practice in bent pipe satellite systems.

Hub terminals periodically broadcast a table which indicates the correspondence between IP multicast address and MPEG PID which carries them. Remote terminals locally store the table, and use it when needed.

With a configured correspondence, a network Manager shall configure the correspondence table in each remote terminals of its network.

10.3.4 Multicast routing

In the case of dynamic multicast groups, a signalling multicast protocol is required on the BSM network to route IP multicast packets. Two kinds of protocols could be suggested depending on BSM scenarios. These two types of protocols required a bi-directional link between Satellite Terminals. The return link could be either a satellite link or a terrestrial link based on the Link Layer Tunnelling Mechanism (IETF RFC 3077 [30]).

For the access network scenario, a group membership protocol is recommended. IGMPv2 [8] is the most widely used protocol on IPv4-enabled local network. A host wishing to receive or to stop receiving IP multicast group data, sends group subscription or de-subscription messages to an elected router in charge of managing group membership. On its side, the elected router will query hosts, in order to check if they are still members on the interface. A single report from host is then required to confirm the checking. In terrestrial networks, there is a mechanism that prevents all hosts from reporting: any host that hears another device confirm the membership should cancel any pending membership replies. However, this mechanism does not work properly in a satellite context, because it requires a shared medium in order to enable the STs to listen to each other. Thus, in a satellite context, all STs will respond to queries and will flood the queries.

For IPv6-enabled local network the IETF is defining a new protocol, MLD (Multicast Listener Discovery) [36]. At present MLD is based on the same mechanism as IGMPv3 (source filtering). A host sends a (source, group) subscription message to an elected router instead of a group subscription message.

These protocols build a membership table containing for each group or (source, group) the list of output interfaces to forward IP multicast packets. This last point is important: with this table a router knows that a group or a (source, group) has at least one subscriber on an interface, it does not know the list of subscribers neither per group (source, group) nor per interface.

For core network and distribution network scenarios, a routing protocol is recommended. The most relevant routing protocol is PIM-SM (IETF RFC 2362 [31]; or PIM-SM v2-revised [34]). This protocol could be used intra-domain and inter-domains. It belongs to the Sparse Mode category: explicit subscription messages are required to build the delivery tree.

Specific configuration of PIM-SM may be required to be efficient on BSM network. A BSM network has a long delay compared to terrestrial network, and a large number of routers listening to the satellite link. An important issue is the location of the Rendez-vous Point (RP) on a BSM network. In terrestrial network some routers are candidate to be RP and a mechanism elects a RP per group. In BSM network the RP location has to be judicious to prevent multicast data (sent by the source to the group RP encapsulated into unicast IP packet) from being forwarded several times on the air interface. A RP could provide RP functions for several groups at the same time. A second issue, linked to the first one, is the switching from the RP-rooted tree to the Source-rooted tree (Shortest Path Tree). Other issues are Hello message sending rate, the values of timers linked to Join/ Prune messages and states.

The PIM-SM protocol uses the unicast routing table to compute multicast delivery trees however it does not compute its own unicast routing table. A static unicast routing table or a routing table managed by a specific protocol is required. To forward IP multicast packets, this protocol builds a multicast routing table indicating for a group or a (source, group) the input interface (also called the Reverse Path Forwarding interface) and the list of the output interfaces. A multicast routing table does not give the next hop router (as a unicast routing table does).

11 Security

11.1 Introduction

11.1.1 IP threats

Different attacks can be applied to perturb an IP session; however, we can consider three big families of IP threats:

a) **IP sniffing:** the intruder listens to the traffic transiting over the network. He can have knowledge of the data included in the messages on the network, as well as users' passwords. IP sniffing is rather easy among broadcast environment.

- The reply to this threat is data confidentiality.

b) **IP spoofing:** the intruder takes the identity of a user. The aim of this attack is to establish a connection with another user or a web site taking the identity of someone else or insert data in a running communication.

- This kind of attack can be done either by modifying the MAC address of the intruder's device, or by creation of ICMP messages in order to redirect IP messages towards the intruder's device, or inclusion of TCP packets with appropriate sequence numbers in a connection.
- The reply to this threat is entity authentication, entity authorization and accountability; and data integrity.

c) **IP flooding:** the intruder sends a vast number of packets to a single entity, provoking a congestion that impeaches the destination to process the legal IP packets.

- The reply to this threat is entity authorization and accountability and availability.
- At this time, there is no guaranteed reply to this intrusion.

11.1.2 BSM security processes

The objective of the BSM network is to provide a wide range of multimedia services for business and consumers. Because of the rich content of data transiting on the network, security becomes a mandatory requirement.

Security is intended to protect the user identity, the signalling traffic and the data traffic. The following security components should be considered for a BSM system:

- **Data Confidentiality:** Confidentiality of stored and transferred information. This means protecting the signalling and/or traffic data. It prevents eavesdroppers from listening to the traffic transiting over the BSM network (sniffing attack, etc.).
- **Data Integrity:** Protection of stored and transferred information against corruption or loss. This means detecting unauthorized modifications to data. It prevents the intruder from modifying data transiting over the BSM network.
- **Entity Authentication:** Verifying and ensuring the identities of the partners involved in establishing the communication. It prevents the intruder from taking the identity of a legitimate user (spoofing, masquerade attacks, etc.).
- **Entity Authorization and Accountability:** Ensuring that any entity should be responsible for any actions initiated. It means identifying and validating the initiating entity for all network service invocations and for all network management activities.
- **Availability:** Ensuring that all legitimate entities should experience correct access to services and facilities of the BSM network. It prevents the intruder from disturbing or misusing the network services leading to a denial of service (flooding, man-in-the-middle attack, etc.).

11.2 Security requirements

11.2.1 General performance requirements

Due to the intrinsic broadcast capability of the BSM satellite system, the radio segment constitutes a large network causing some critical security concerns.

The BSM security design shall focus on internal interfaces, in particular the air interface between the hub and a satellite terminal in the case of a star network architecture; or between satellite terminals in the case of a mesh network architecture.

The security design shall also take account of the features of a satellite system such as the delay and the scarce bandwidth. The following general performance requirements shall be applied to the security system:

- Overhead reduction in order to limit the waste of bandwidth;
- Number of signalling messages reduction (authentication, keys exchange) between the stations in order to limit the latency for the connection setup and the overhead.

11.2.2 Compatibility requirements

Analyzing the role and the interest of the different actors (Access Network Operator, Internet Service Provider), it can be possible to have simultaneously different security scheme:

- On the BSM network between the hub and the terminals for the star network case.
- On the BSM network between 2 or more satellite terminals for the mesh network case.
- On the end-to-end communication between the Service Provider servers and the CPE client.

The security functions on the BSM system shall be compatible with an end-to-end security architecture that may be provided by a Service Provider.

The security functions shall also be compatible with some widely used networking functions such as NAT (Network Address Translation) and PEPs (Performance Enhancing Proxies) such as TCP acceleration.

11.2.3 Services requirements

The BSM system should implement some native mechanisms to provide secure multicast transmission of IP packets on the satellite bi-directional links. The key issues are to implement a secure access control, a scalable key management and an efficient ciphered data-plane which are optimized for multicast environment. Usual key management and entity authentication schemes are defined for point-to-point or pure broadcast communications.

In point to point schemes, cryptographic functions and keys are negotiated, and both entities authenticate each other. These schemes are not suitable for multicast communications.

In pure broadcast systems without return link, the key distribution is often based on pre-shared keys.

The BSM system requires a specific key management system and a secure access control, optimized for IP multicasting. The therefore most optimized scheme would be based on a centralized management, achieved by a central server, which would manage multicast groups, authenticate ST at their entry in the network, and distribute securely group session keys. (the group session keys are used to encrypt data transmitted on satellite links and destined to multicast groups).

The other requirements of such a mechanism would be:

- To support unicast and multicast communications.
- To allow every ST to be receiver and sender.
- To be compatible with any data ciphering protocol.
- To be able to manage simultaneously multiple multicast groups and to ensure the security of each one.

- To provide specific source data authentication mechanisms, i.e. to ensure that the traffic comes from another group member (group authentication) or from a particular entity (individual authentication).

11.3 Security in the satellite independent layers

11.3.1 General

Several different possibilities exist to secure the traffic:

- At the application level providing:
  - Data confidentiality using SSL (Secure Socket Level), PGP (Pretty Good Privacy), TLS (Transport Layer Security) or others.
  - Entity authentication using PAP, CHAP or others.
- At the network level providing:
  - Authentication and data confidentiality using IPSec.

Security techniques at the satellite independent layers include the techniques used by Service Providers for end-to-end security.

11.3.2 IPSec

The goal of IPSec is to provide security at the IP layer. IPSec defines three separate components:

- Authentication Header (AH): defined in IETF RFC 2402 [27], provides connectionless integrity services, authentication, anti-replay.
- Encapsulated Security Payload (ESP): defined in IETF RFC 2406 [28], provides data confidentiality by encryption, flow confidentiality control, connectionless integrity, authentication, anti-replay.
- Internet Key Exchange (IKE): defined in IETF RFC 2409 [29].

A unidirectional connection, called Security Association (SA) provides security services by implementing AH or ESP. An SA is identified by:

- A Security Parameter Index (SPI).
- The destination IP address.
- A security protocol identifier (AH or ESP).

The existence of two different protocols (AH and ESP) is justified by the fact that if only authentication is needed, AH is sufficient as it is not regulated and thus can be used freely. In another hand, if authentication is needed for sensitive configurations, some local rules regarding encryption are required and implementation of ESP is necessary.

IPSec can be configured end-to-end or over a tunnel (gateway-to-gateway).

IPSec is optional in IPv4, but is mandatory in IPv6.

In IPv6, IPSec is located just after the IP header and its extensions, i.e. Hop-by-Hop, Routing and Fragmentation headers, but before any possible destination options.

11.3.3 Constraints on the use of IPSec

IPSec has constraints that do not allow it to be used with all kinds of configurations or protocols. Those limitations are:

- IPSec is originally configured for unicast sessions. Multicast is however integrated for IPSec over IPv6.
- IPSec ESP tolerates no modification of the Layer 4 headers:
  - This implies that Header Compression, whose role is to compress IP and TCP or UDP headers in order to reduce the datagrams overhead, cannot be implemented with IPSec ESP, as the TCP/UDP header is modified.
  - This also implies that PEPs (Performance Enhancing Proxies), for example a PEP that is used to spoof the TCP header, cannot not be used with IPSec ESP.
- ESP has different levels of encryption. The highest levels are or may be subject to local authorities (i.e. governments and/or regulation authorities) control and deliverance. These levels may change from a country to another (length tolerance of the keys):
  - In a Worldwide environment, it is important that the security level is the same whatever the user's or gateway location. This implies agreement among standardization bodies and authorities on the usage of ESP.
- In the tunnel mode, IPSec is implemented between two gateways that may be located in different premises than the users'. It is then assumed that the IPSec configuration is static:
  - It is however conceivable that commands could allow IPSec to be remote controlled.

Therefore the IPSec technique leads to some critical issues in a satellite context:

- End-to-end IPSec can not run simultaneously with techniques such as PEPs and NAT.
- IPSec can not support multicast (because of an unicast key exchange, IKE):
  - However, the security issues for multicast services are currently under study in the IETF MSEC (multicast security) working group [35].
- IPSec requires that the same encryption rules must be applied on both ends.
- The most implemented IPSec mode is the Tunnel mode leading to a critical overhead and a large number of signalling messages (high latency during setup).

11.4 Security in the satellite dependent layers

11.4.1 General

Security techniques implemented in the satellite dependent layers enables:

- to focus on the BSM internal interfaces (as defined in clause 11.2.1);
- to offer a total compatibility with higher layer security systems.

11.4.2 DVB techniques

At the data link level, DVB defines several different techniques but these do not provide a full suitability for BSM:

- DVB-CA is available for DVB-S compatible systems to provide scrambling at the level of the Transport Stream packet (MPEG2 184 bytes packets). This technique is already deployed for digital television and is also used for IP data confidentiality.

- DVB-RCS is defining a security mechanism at the data link layer for the individual user scrambling. Three security mechanisms are proposed:
  - Authentication of the ST to the NCC and key-agreement between the NCC and ST during a session set-up with a set of MAC messages (also used for key updating).
  - Defence against clones.
  - Encryption/ Decryption of payload data streams between the NCC and the ST.

The DVB-RCS security mechanisms may be considered as an interesting basis for BSM security in the Satellite Dependent Layers but they also requires some adaptations such as:

- The expansion of the services to data traffic between the Hub terminal and the Satellite terminals (should run without any modification).
- An additional mechanism to provide a full multicast support.
- The encryption applied on data and/or signalling traffic.
- The possible use of smartcards.

12 Performance Enhancing Proxies (PEPs)

12.1 Overview of PEPs

This clause discusses Performance Enhancing Proxy (PEP) performance mitigation techniques as applied to a BSM system. The discussion of PEPs in this clause is based on the discussion in IETF RFC 3135 [20].

A PEP can be used to improve the performance of the Internet protocols on network paths where native performance suffers due to characteristics of a link or subnetwork on the path. PEPs are typically focussed on improving throughput for applications such as FTP and HTTP web page retrievals. To a lesser degree, PEP implementations also work to improve interactive response time for small transactions.

12.2 Definitions

12.2.1 Layering

A PEP implementation may function at any protocol layer but typically it functions at one or two layers only.

Consideration of PEP implementations for BSM will focus on implementations that function at the transport layer or at the application layer as such PEPs are most commonly used to enhance performance over links with problematic characteristics.

NOTE: A PEP implementation may also operate below the network layer, that is, at the link layer, but such link layer mechanisms can be and typically are implemented transparently to network and higher layers, requiring no modifications to protocol operation above the link layer.

12.2.2 Implementation distribution

A PEP implementation may be integrated, i.e. it comprises a single PEP component implemented within a single node, or distributed, i.e. it comprises two or more PEP components, typically implemented in multiple nodes.

For example, a distributed PEP implementation is most appropriate for a satellite link for which performance enhancement is desired. A typical implementation would comprise two PEPs located at each end of the satellite link.

12.2.3 Implementation symmetry

A PEP implementation may be symmetric or asymmetric. Symmetric PEPs use identical behaviour in both directions, i.e. the actions taken by the PEP occur independent from which interface a packet is received. Asymmetric PEPs operate differently in each direction.

The direction can be defined in terms of the link (e.g. from a hub station to a remote station) or in terms of protocol traffic (e.g. the direction of TCP data flow, often called the TCP data channel, or the direction of TCP ACK flow, often called the TCP ACK channel). For example, an asymmetric application layer PEP might be used for the request-reply type of HTTP traffic.

A PEP implementation may also be both symmetric and asymmetric at the same time with regard to different mechanisms it employs (PEP mechanisms are described in IETF RFC 3135 [20], clause 3).

Whether a PEP implementation is symmetric or asymmetric is independent of whether the PEP implementation is integrated or distributed. In other words, a distributed PEP implementation might operate symmetrically at each end of a link (i.e. the two PEPs function identically). On the other hand, a distributed PEP implementation might operate asymmetrically, with a different PEP implementation at each end of the link.

12.2.4 Split TCP connections

A split connection TCP implementation terminates the TCP connection received from an end system and establishes a corresponding TCP connection to the other end system.

In a distributed PEP implementation, this is typically done to allow the use of a third connection between two PEPs optimized for the satellite link. This third connection might be a TCP connection optimized for the link or it might be another protocol, for example, a proprietary protocol running on top of UDP. Also, the distributed implementation might use a separate connection between the proxies for each TCP connection or it might multiplex the data from multiple TCP connections across a single connection between the PEPs.

In an integrated PEP split connection TCP implementation, the PEP again terminates the connection from one end system and originates a separate connection to the other end system.

12.2.5 Transparency

A transparent PEPs is one that operates totally transparently to the end systems, transport endpoints, and/or applications involved (in a connection), requiring no modifications to the end systems, transport endpoints, or applications.

A non-transparent PEP implementation is one that may require modifications to one or both ends in order to be used. Both of these kinds of PEP implementations are non-transparent, at least to the layer requiring modification.

It is sometimes useful to think of the degree of transparency of a PEP implementation at four levels:

- transparency with respect to the end systems (network-layer transparent PEP);
- transparency with respect to the transport endpoints (transport-layer transparent PEP);
- transparency with respect to the applications (application-layer transparent PEP);
- transparency with respect to the users.

For example, a user who subscribes to a satellite Internet access service may be aware that the satellite terminal is providing a performance enhancing service even though the TCP/IP stack and the applications in the user's PC are not aware of the PEP which implements it.

Note that the issue of transparency is not the same as the issue of maintaining end-to-end semantics. For example, a PEP implementation which simply uses a TCP ACK spacing mechanism maintains the end-to-end semantics of the TCP connection while a split connection TCP PEP implementation may not.

12.3 Negative implications of using PEPs

12.3.1 General

Most of the potential negative implications associated with using PEPs are related to the possibility of breaking the end-to-end semantics of connections. This is one of the main reasons why PEPs are not recommended for general use.

NOTE: The end-to-end argument is one of the architectural principles of the Internet. The basic argument is that, as a first principle, certain required end-to-end functions can only be correctly performed by the end systems themselves.

As indicated in IETF RFC 3135 [20] not all PEP implementations break the end-to-end semantics of connections. Correctly designed PEPs do not attempt to replace any application level end-to-end function, but only attempt to add performance optimizations to a subpath of the end-to-end path between the application endpoints.

12.3.2 Security implications

In most cases, security applied above the transport layer can be used with PEPs, especially transport layer PEPs.

The most detrimental negative implication of breaking the end-to-end semantics of a connection is that it disables the use of IPSEC. In general, a user or network administrator must choose between using PEPs and using IPSEC.

If IPSEC is employed end-to-end, PEPs that are implemented on intermediate nodes in the network cannot examine the transport or application headers of IP packets because encryption of IP packets via the IPsec ESP header (in either transport or tunnel mode) renders the TCP header and payload unintelligible to the PEPs. Without being able to examine the transport or application headers, a PEP may not function optimally or at all.

The same limitation applies between IPSEC firewalls at user-specified intermediate points This applies to the use of IPSEC in tunnel mode, when the tunnel end-points are not under the control of the satellite operator.

12.3.3 Fate sharing

If the end-to-end connection depends upon some state being stored in the network (e.g. in a PEP), then a failure in the network (e.g. the node containing a PEP crashes) causes this state to be lost, forcing the connection to terminate even if an alternate path through the network exists.

In the case of BSM networks PEPs are typically used in an environment where there is no alternate path between the end systems and, therefore, a failure of the intermediate node containing a PEP would result in the termination of the connection in any case.

Accepting this risk should be under the control of the user (i.e. the user should always have the option to choose end-to-end operation) and, if the user chooses to use the PEP, the user should be aware of the implications that a PEP failure has with respect to the applications being used.

12.3.4 End-to-end reliability

If a PEP implementation acknowledges application data prematurely (before the PEP receives an application ACK from the other endpoint), end-to-end reliability cannot be guaranteed. Typically, application layer PEPs do not acknowledge data prematurely, i.e. the PEP does not send an application ACK to the sender until it receives an application ACK from the receiver. And, transport layer PEP implementations, including TCP PEPs, generally do not interfere with end-to-end application layer acknowledgements as they let applications operate end-to-end.

However, the user and/or network administrator employing the PEP must understand how it operates in order to understand the risks related to end-to-end reliability.

12.4 Alternatives to using PEPs

12.4.1 Alternative transport layer protocols

An alternative to the use of PEPs, is to use an alternative end-to-end transport protocol in place of TCP. This approach should preserve the end-to-end behaviour of the transport layer (see below for the implications) while still permitting part of the protocol to be matched to the characteristics of a link or subnetwork on the network path.

An example of the basic principle is illustrated in figure 12.4.1. In this example, the transport layer is divided into two sub-layers: the upper sublayer is responsible for maintaining the end-to-end reliability; and the lower sublayer is responsible for the flow control over the satellite portion of the network path (i.e. the BSM subnetwork). The lower sublayer - the Satellite Transport Layer - provides an optimized transport layer for the satellite portion of the path.

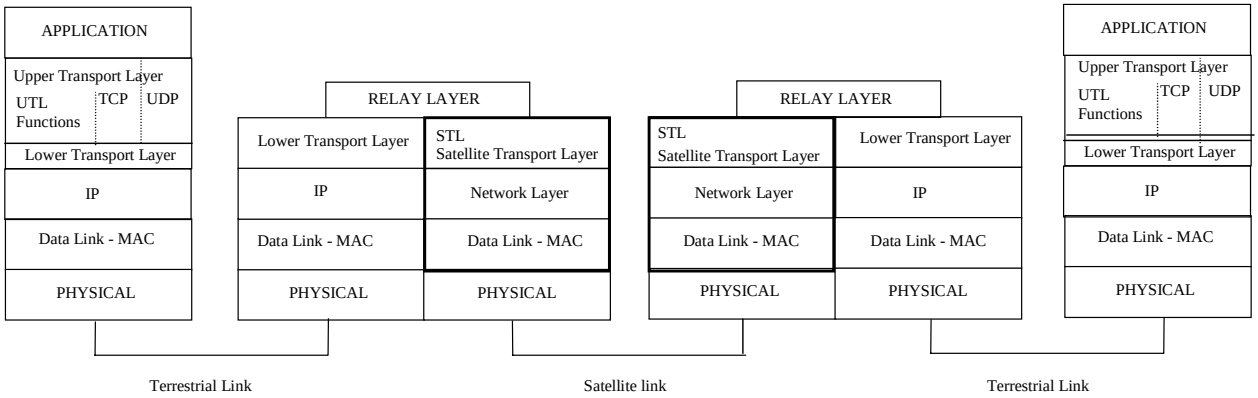

Figure 12.4.1: Example of modified Transport Layer protocol

This approach has implications that include:

a) The end-to-end behaviour of the transport protocol should be preserved. In particular, the end-to-end reliability between the communicating end hosts should be guaranteed.

b) A requirement that both of the end hosts implement the alternative transport layer protocol. In order to maintain full interoperability with a range of different end hosts, this approach must therefore include the ability to revert to using the standard TCP transport protocol in the event that one of the end hosts does not support the alternative protocol.

c) There is a requirement to provide fairness between flows that share a common part of an Internet path. This implies that any new protocol must adopt a congestion control procedure that is accepted for use within the general Internet.

d) There are architectural issues introduced when there are alternate transport protocols, e.g. how an application/host chooses which protocol to use, when presented with a set of possible transport protocols offering different feature sets (e.g. ECN capability (see [26], mobility extensions, satellite-specific enhancements).

There are therefore important architectural and protocol issues to be considered before deploying or modifying transport protocols for the TCP/IP protocol suite. The IETF has not currently recommended a subnetwork-specific transport protocol for use in the general Internet.

12.4.2 Space Communications Protocol Standards (SCPS)

The goal of the Space Communications Protocol Standards (SCPS) project [45] is to provide a suite of standard data handling protocols that (from a user viewpoint) make a remote space vehicle appear to be just another "node on the Internet".

A protocol of interest for PEPs is the SCPS Transport Protocol (SCPS-TP) [46], [47] which is optimized to provide reliable end-to-end delivery of spacecraft command and telemetry messages between computers that are communicating over a network containing one or more potentially unreliable space data transmission paths. SCPS-TP is based on TCP and UDP with modifications to suit the spacecraft communications environment. Depending on configuration options selected, SCPS-TP can be less than TCP, identical to TCP, or a super-set of TCP.

Annex A:
Segmentation and fragmentation

The definitions used in the present specification are illustrated in figure A.1.

The terms Protocol Data Unit (PDU) and Service Data Unit (SDU) are used in accordance with the ISO Open Systems Interconnection Basic Reference Model [6].

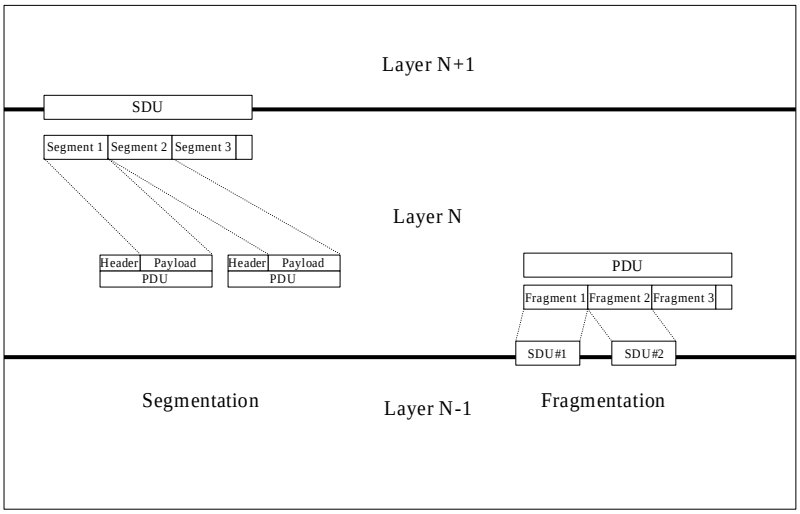

Figure A.1: Segmentation and fragmentation

Segmentation is the process whereby a single Service Data Unit (SDU) is split into one or more Protocol Data Units (PDUs) for transmission to a peer entity.

Fragmentation is the process whereby a single Protocol Data Unit (PDU) is split into one or more Service Data Units (SDUs) for submission to the next lower layer.

As illustrated in figure A.1, a given layer of protocol may contain both segmentation and fragmentation. Two examples are noted:

- TCP supports segmentation of user SDUs; e.g. where the user data stream is segmented into a series of TCP datagrams;
- IP supports fragmentation of IP packets (IP PDUs) where a large IP packet may be split into a series of smaller IP packets.

Annex B: Bibliography

- Venkata N. Padmanabhan, Hari Balakrishnan, Gorry Fairhurst, Mahesh Sooriyabandara "TCP Performance Implications of Network Asymmetry" draft-ietf-pilc-asym-04.txt, Work In Progress, IETF PILC WG.
- IETF RFC 791: "Internet Protocol".
- IETF RFC 2460: "Internet Protocol, Version 6 (IPv6) Specification".

History

Document history		
V1.1.1	November 2002	Publication