# ETSI TR 101 545-5 V1.1.1 (2014-04)


**Digital Video Broadcasting (DVB);**
**Second Generation DVB Interactive Satellite System (DVB-RCS2);**
**Part 5: Guidelines for the Implementation and Use of TS 101 545-3**

EBU
OPERATING EUROVISION

DVB
Digital Video Broadcasting

Reference
DTR/JTC-DVB-324-5

Keywords
DVB, interaction, satellite

***ETSI***

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

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

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

# Contents

# Intellectual Property Rights

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

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 Joint Technical Committee (JTC) Broadcast of the European Broadcasting Union (EBU), Comité Européen de Normalisation ELECtrotechnique (CENELEC) and the European Telecommunications Standards Institute (ETSI).

NOTE: The EBU/ETSI JTC Broadcast was established in 1990 to co-ordinate the drafting of standards in the specific field of broadcasting and related fields. Since 1995 the JTC Broadcast became a tripartite body by including in the Memorandum of Understanding also CENELEC, which is responsible for the standardization of radio and television receivers. The EBU is a professional association of broadcasting organizations whose work includes the co-ordination of its members' activities in the technical, legal, programme-making and programme-exchange domains. The EBU has active members in about 60 countries in the European broadcasting area; its headquarters is in Geneva.

European Broadcasting Union
CH-1218 GRAND SACONNEX (Geneva)
Switzerland
Tel: +41 22 717 21 11
Fax: +41 22 717 24 81

The Digital Video Broadcasting Project (DVB) is an industry-led consortium of broadcasters, manufacturers, network operators, software developers, regulatory bodies, content owners and others committed to designing global standards for the delivery of digital television and data services. DVB fosters market driven solutions that meet the needs and economic circumstances of broadcast industry stakeholders and consumers. DVB standards cover all aspects of digital television from transmission through interfacing, conditional access and interactivity for digital video, audio and data. The consortium came together in 1993 to provide global standardization, interoperability and future proof specifications.

The present document is part 5 of a multi-part deliverable covering the DVB Interactive Satellite System specification as identified below:

TS 101 545-1: "Overview and System Level specification";

EN 301 545-2: "Lower Layers for Satellite standard";

TS 101 545-3: "Higher Layers Satellite Specification";

TR 101 545-4: "Guidelines for Implementation and Use of EN 301 545-2";

**TR 101 545-5: "Guidelines for the Implementation and Use of TS 101 545-3".**

# Introduction

TS 101 545-3 [i.1] provides the specification of the higher-layer satellite architecture, signalling, and functions required for the two way interactive satellite networks that are specified in [i.2], and [i.3] together with its implementation guidelines [i.4]. The requirements in [i.1] have been introduced to provide the best possible interoperability between terminals and hubs, defining the network functions as well as management and control capabilities to complement the lower-layer specification of the system (up to and including layer 2) given in [i.3].

The present document provides guidelines for the implementation and the usage of the higher-layer architectural elements and functions that are described in [i.1]. It is aimed that the present document completes [i.1] with implementation and configuration examples, recommended practices, and informative elaborations to help attain full terminal-hub interoperability as far as higher-layer functionalities are concerned. The present document often refers to MIB objects that are defined in [i.1] and lower-layer signalling tables/descriptors that are defined in [i.3]. In addition to [i.1] and [i.3], [i.4] contains lower-layer descriptions and recommendations that are useful to complement the discussions in the present document. This is particularly the case in the discussions on QoS support and satellite virtual networks.

The present document covers transparent star, regenerative mesh, and transparent mesh overlay network topologies.

Clause 2 provides the references. Clause 3 provides the definitions, explains symbols, and expands abbreviations.

Clause 4 provides further guidance in the reading of the present document through the introduction of reference models.

Clause 5 elaborates on IP routing support over the satellite interface, and provides guidance on dynamic routing support using OSPF. The clause also provides recommendations on the usage of OSPF in mesh satellite networks.

Clause 6 provides recommendations on IP multicast support in transparent star network topology. The clause also provides L3/L2 address mapping examples using DVB-RCS2 lower-layer signalling.

Clause 7 elaborates on the QoS model defined for DVB-RCS2. The clause refers to cardinality diagram describing the relationship among different user and control plane entities. It also provides configuration examples and recommendations for transparent star and regenerative mesh networks. QoS support in DVB-RCS2 relies on lower-layer service description and entities. The present document scope is confined to higher layers with minimal overlaps. The reader is referred to [i.3] and [i.4] for lower-layer QoS support.

Clause 8 elaborates on satellite virtual networking and virtual LANs. The clause provides recommendations on VLAN support. It also provides recommendations on mapping SVN tags to return link encapsulation fields.

Clause 9 provides an example PEP negotiation protocol that uses PEP messages that are defined in [i.1].

Clause 10 provides recommendations on default and operational SNMP configuration for the different management actors/roles in the network.

Clause 11 provides a step-by-step walkthrough of the startup procedure in DVB-RCS2 terminals.

Clause 12 provides an example of OSS-NMC interface that is aligned with 3GPP specifications. The clause also elaborates on performance management, and provides an example list of key performance indicators.

Clause 13 elaborates on Dynamic Connectivity Protocol for regenerative mesh networks. The clause also provides example state-transition diagrams and message sequence diagrams.

Clause 14 provides an encompassing description of transparent mesh overlay network support in DVB-RCS2. The clause includes elaborations on routing, multicasting, QoS support; specifically in transparent mesh overlay networks.

Clause 15 provides guidance and example message sequence diagrams on Dynamic Connectivity Protocol for transparent mesh overlay networks.

Annex A provides recommendations in regards to the integration of DVB-RCS2 interactive networks with the service layer of the Next Generation Networks (NGN) architecture. Annex B provides recommendation and guidelines for efficient deployment of secure VPNs in broadband satellite systems. Annex C elaborates on and provides recommendations for TCP transport protocol in the presence of random access user data transmission on the satellite return link.

# 1 Scope

The present document provides implementation and usage guidelines for higher-layer functions in DVB-RCS2 interactive satellite networks, which is defined in [i.2]. The lower-layer specification and implementation guidelines for DVB-RCS2 networks are presented in [i.3] and [i.4], respectively.

The present document covers on transparent star, regenerative mesh, and transparent mesh overlay network topologies. The recommendations and examples provided in the present document are informative.

# 2 References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

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

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

## 2.1 Normative references

Not applicable.

## 2.2 Informative references

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] ETSI TS 101 545-3: "Digital Video Broadcasting (DVB); Second Generation DVB Interactive Satellite System (DVB-RCS2); Part 3: Higher Layers Satellite Specification".

[i.2] ETSI TS 101 545-1: "Digital Video Broadcasting (DVB); Second Generation DVB Interactive Satellite System (DVB-RCS2); Part 1: Overview and System Level specification".

[i.3] ETSI EN 301 545-2 V1.1.1 (2012-01): "Digital Video Broadcasting (DVB); Second Generation DVB Interactive Satellite System (DVB-RCS2); Part 2: Lower Layers for Satellite standard".

[i.4] ETSI TR 101 545-4: "Digital Video Broadcasting (DVB); Second Generation DVB Interactive Satellite System (DVB-RCS2); Part 4: Guidelines for Implementation and Use of EN 301 545-2".

[i.5] IETF RFC 2328: "OSPF Version 2".

[i.6] IETF RFC 2453: "RIP Version 2".

[i.7] IETF RFC 5340: "OSPF for IPv6".

[i.8] IETF RFC 4271: "A Border Gateway Protocol 4 (BGP-4)".

[i.9] IETF RFC 5880: "Bidirectional Forwarding Detection (BFD)".

[i.10] IETF RFC 5881: "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)".

[i.11] IETF RFC 1112: "Host Extensions for IP Multicasting".

[i.12] IETF RFC 2365: "Administratively Scoped IP Multicast".

[i.13] IETF RFC 2236: "Internet Group Management Protocol, Version 2".

[i.14] IETF RFC 3376: "Internet Group Management Protocol, Version 3".

[i.15] IETF RFC 4606: "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control".

[i.16] IETF RFC 3810: "Multicast Listener Discovery Version 2 (MLDv2) for IPv6".

[i.17] IETF RFC 4601: "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)".

[i.18] IETF RFC 4605: "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")".

[i.19] IETF RFC 4541: "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches".

[i.20] IETF RFC 3171: "IANA Guidelines for IPv4 Multicast Address Assignments".

[i.21] IETF RFC 2474: "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers".

[i.22] IETF RFC 2475: "An Architecture for Differentiated Services".

[i.23] IETF RFC 3290: "An Informal Management Model for Diffserv Routers".

[i.24] IETF RFC 3086: "Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification".

[i.25] IETF RFC 2753: "A Framework for Policy-based Admission Control".

[i.26] IETF RFC 2698: "A Two Rate Three Color Marker".

[i.27] IETF RFC 2697: "A Single Rate Three Color Marker".

[i.28] IETF RFC 3246: "An Expedited Forwarding PHB (Per-Hop Behavior)".

[i.29] IETF RFC 3247: "Supplemental Information for the New Definition of the EF PHB (Expedited Forwarding Per-Hop Behavior)".

[i.30] IETF RFC 2597: "Assured Forwarding PHB Group".

[i.31] IETF RFC 4594: "Configuration Guidelines for DiffServ Service Classes".

[i.32] IETF RFC 3584: "Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework".

[i.33] IETF RFC 3413: "Simple Network Management Protocol (SNMP) Applications".

[i.34] IETF RFC 3415: "View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)".

[i.35] IETF RFC 3411: "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks".

[i.36] IETF RFC 3918: "Methodology for IP Multicast Benchmarking".

[i.37] IETF RFC 3412: "Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)".

[i.38] IETF RFC 3414: "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)".

[i.39] IETF RFC 5728: "The SatLabs Group DVB-RCS MIB".

[i.40] ETSI TS 132 101: "Digital cellular telecommunications system (Phase 2+); Universal Mobile Telecommunications System (UMTS); LTE; Telecommunication management; Principles and high level requirements (3GPP TS 32.101)".

[i.41] ETSI TS 132 150: "Digital cellular telecommunications system (Phase 2+); Universal Mobile Telecommunications System (UMTS); LTE; Telecommunication management; Integration Reference Point (IRP) Concept and definitions (3GPP TS 32.150)".

[i.42] 3GPP TS 42.435: "Telecommunication management; Performance measurement; eXtensible Markup Language (XML) file format definition".

[i.43] ETSI TS 132 300: "Digital cellular telecommunications system (Phase 2+); Universal Mobile Telecommunications System (UMTS); LTE; Telecommunication management; Configuration Management (CM); Name convention for Managed Objects (3GPP TS 32.300)".

[i.44] ETSI TS 132 405: "Digital cellular telecommunications system (Phase 2+); Universal Mobile Telecommunications System (UMTS); LTE; Telecommunication management; Performance Management (PM); Performance measurements; Universal Terrestrial Radio Access Network (UTRAN) (3GPP TS 32.405)".

[i.45] 3GPP TS 22.228 V12.0.0 (2011-12): "Service requirements for the Internet Protocol (IP) Multimedia core network Subsystem (IMS), Stage 1".

[i.46] 3GPP TS 23.203 V11.4.0 (2011-12): "Technical Specification Group Services and System Aspects; Policy and charging control architecture (Release 11)".

[i.47] ETSI TS 132 240: "Digital cellular telecommunications system (Phase 2+); Universal Mobile Telecommunications System (UMTS); LTE; Telecommunication management; Charging management; Charging architecture and principles (3GPP TS 32.240)".

[i.48] 3GPP TS 23.402 V11.0.0 (2011-09): "Technical Specification 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Architecture enhancements for non-3GPP accesses (Release 11)".

[i.49] 3GPP TS 29.212 V11.2.0 (2011-09): "Technical Specification 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Policy and Charging Control (PCC) over Gx/Sd reference point (Release 11)".

[i.50] 3GPP TS 29.213 V11.0.0 (2011-09): "Technical Specification 3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Policy and Charging Control signalling flows and Quality of Service (QoS) parameter mapping (Release 11)".

[i.51] IETF RFC 3588: "Diameter Base Protocol".

[i.52] IETF RFC 5213: "Proxy Mobile IPv6".

[i.53] IETF RFC 4301: "Security Architecture for the Internet Protocol".

[i.54] IETF RFC 3522: "The Eifel Detection Algorithm for TCP".

[i.55] IETF RFC 4015: "The Eifel Response Algorithm for TCP".

[i.56] IETF RFC 5682: "Forward RTO-Recovery (F-RTO): An Algorithm for Detecting Spurious Retransmission Timeouts with TCP".

[i.57] M. Fiedler, T. Hossfeld, and P. Tran-Gia: "A Generic Quantitative Relationship between Quality of Experience and Quality of Service" IEEE Network, vol. 24, no. 2, Apr. 2010, pp. 36-41.

[i.58] H. Skinnemoen, A. Vermesan, A. Iuoras, G. Adams, and X. Lobao: "VoIP over DVB-RCS with QoS and bandwidth on demand" IEEE Wireless Communications, vol.12, no.5, pp. 46- 53, Oct. 2005.

[i.59] IETF RFC 5681: "TCP Congestion Control".

[i.60] IETF RFC 6298: "Computing TCP's Retransmission Timer".

[i.61] N. Dukkipati, T. Refice, Y. Cheng, J. Chu, N. Sutin, A. Agarwal, T. Herbert, and J. Arvind: "An Argument for Increasing TCP's Initial Congestion Window", ACM SIGCOMM Computer Communications Review, vol. 40, pp. 27-33, July 2010.

[i.62] J. Chu, N. Dukkipati, Y. Cheng, and M. Mathis: "Increasing TCP initial Window", Internet Draft, draft-hkchu-tcpm-initcwnd-01.txt, July 2010.

[i.63] D. J. Wischik: "Short Messages", Philosophical Transactions of the Royal Society A, vol. 366, pp. 1941-1953, 2008.

[i.64] IETF RFC 5690: "Adding Acknowledgement Congestion Control to TCP".

[i.65] IETF RFC 3390: "Increasing TCP's Initial Window".

[i.66] Y. Chen: "Seeding RTO with RTT sampled during three-way handshake", Internet Draft, draft-ycheng-tcpm-rtosynrtt.txt, IETF, June 2010.

[i.67] Y. Chen, J. Chu, and A. Jain: "TCP Fast Open", Internet Draft, draft-cheng-tcpm-fastopen-00.txt, March 2011.

[i.68] IETF RFC 793: "Transmission Control Protocol".

[i.69] IETF RFC 1901: "Introduction to Community-based SNMPv2".

[i.70] IETF RFC 1905: "Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)".

[i.71] IETF RFC 5882: "Generic Application of Bidirectional Forwarding Detection (BFD)".

[i.72] IEEE 802.1pQ: "IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks".

# 3 Definitions, symbols and abbreviations

## 3.1 Definitions

For the purposes of the present document, the terms and definitions given in [i.1] apply.

## 3.2 Symbols

For the purposes of the present document, the symbols given in [i.1] apply.

## 3.3 Abbreviations

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

| | |
|---|---|
| 3GPP | 3rd Generation Project Partnership |
| 3WHS | 3-Way HandShake |
| AAA | Authentication Authorization Accounting |
| AAR | Authentication Authorisation Request |
| ABR | Area Border Router |
| AC | Allocation Channel |
| ACK | ACKnowledgement |
| ACM | Adaptive Coding and Modulation |
| ADC | Application Detection and Control |
| AF | Assured Forwarding |
| AH | Authentication Header |
| ALPDU | Addressed Link Protocol Data Unit |
| AR | Address Resolution |
| ARP | Allocation and Retention Priority |
| AS | Autonomous System |
| ATM | Asynchronous Transfer Mode |
| AVBDC | Absolute Volume Based Dynamic Capacity |

| | |
|---|---|
| AVP | Attribute Value Pair |
| BA | Behaviour Aggregate |
| BBERF | Bearer Binding and Event Reporting Function |
| BCT | Broadcast Configuration Table |
| BDR | Backup Designated Router |
| BE | Best Effort |
| BER | Bit Error Rate |
| BFD | Bidirectional Forwarding Detection |
| BGP | Border Gateway Protocol |
| BSS | Business Support System |
| BW | Bandwidth |
| CA | Connectivity Aggregate |
| CAC | Connection Admission Control |
| CAN | Connectivity Access Network |
| CC | Connectivity Channel or Capacity Category |
| CCA | Credit Control Answer |
| CCR | Credit Control Request |
| CEA | Capabilities Exchange Answer |
| CER | Capabilities Exchange Request |
| CFI | Canonical Format Indicator |
| CLI | Command Line Interface |
| CM | Configuration Management |
| CNG | Customer Network Gateway |
| CNR | Carrier to Noise Ratio |
| COMSEC | Communication Security |
| CoS | Class of Service |
| CPN | Customer Premises Network |
| CPU | Central Processor Unit |
| CR | Capacity Request |
| CRA | Constant Rate Assignment |
| CS | Class Selector |
| CSCF | Call Session Control Function |
| CW | Continuous Wave |
| DA | Dedicated Access |
| DA-AC | Dedicated Access Allocation Channel |
| DAMA | Demand Assignment Multiple Access |
| DCP | Dynamic Connectivity Protocol |
| DF | Default Forwarding |
| DHCP | Dynamic Host Control Protocol |
| DN | Distinguished Name |
| DNS | Domain Name Server |
| DR | Designate Router |
| DS | Differentiated Services |
| DSCP | Differentiated Services Code Point |
| DUP | DUPlicate |
| DVB | Digital Video Broadcast |
| DVB-RCS2 | Digital Video Broadcast Return Channel via Satellite 2nd generation |
| DVB-S | Digital Video Broadcast - Satellite |
| ECN | Explicit Congestion Notification |
| EF | Expedited Forwarding |
| EIRP | Effective Isotropically Radiated Power |
| EM | Elements Manager |
| ESP | Encapsulating Security Payload |
| FCA | Free Capacity Allocation |
| FCAPS | Fault, Configuration, Accounting, Performance, Security |
| FCT2 | Frame Composition Table 2nd generation |
| FEC | Forward Error Correction |
| FIB | Forwarding Information Base |
| FIFO | First In First Out |
| FL | Forward Link |
| FPDU | Frame Protocol Data Unit |

| | |
|---|---|
| FTP | File Transfer Protocol |
| GBR | Guaranteed Bit Rate |
| GPS | Global Positioning System |
| GRE | Generic Routing Encapsulation |
| GSE | Generic Stream Encapsulation |
| GSM | Global System Mobile |
| GW | GateWay |
| HID | Hardware IDentifier |
| HL | Higher Layer |
| HLID | Higher Layers Initialisation Descriptor |
| HLS | Higher Layer Service |
| HSS | Home Subscriber Server |
| HTTP | Hyper-Text Transfer Protocol |
| HTTPS | Hyper-Text Transfer Protocol Secure |
| HW | Hardware |
| IANA | Internet Assigned Numbers Agency |
| IB | Installation Burst |
| ICMP | Internet Control Message Protocol |
| IE | Information Element |
| IETF | Internet Engineering Task Force |
| IF | Intermediate Frequency |
| IFL | Inter-Facility Link |
| IGMP | Internet Group Management Protocol |
| IMS | Internet Multimedia Subsystem |
| IMSI | International Mobile Subscriber Identity |
| INID | Interactive Network ID |
| IP | Internet Protocol |
| IP-CAN | Internet Protocol – Connectivity Access Network |
| IPTV | Internet Protocol TV |
| IPv4 | Internet Protocol version 4 |
| IPv6 | Internet Protocol version 6 |
| IRP | Interface Reference Point |
| ISC | IMS Service Control |
| ISO | International Standards Organisation |
| ISP | Internet Service Provider |
| ITU-T | International Telecommunication Union - Telecommunication |
| IW | Initial Window |
| KPI | Key Performance Indicator |
| L2S | Layer-2 Signalling |
| LAN | Local Area Network |
| LB | Link Behaviour |
| LCD | Local Configuration Datastore |
| LDN | Local Distinguished Name |
| LI | Link Interface |
| LL | Lower Layer |
| LLS | Lower Layer Service |
| LMA | Local Mobility Anchor |
| LNB | Low Noise Block |
| LQC | Link Quality Control |
| LS | Link Stream |
| LSA | Link State Advertisement |
| LSB | Least Significant Bit |
| LSE | Link Service Establishment |
| LTE | Long Term Evolution |
| LW | Loss Window |
| MAC | Medium Access Control |
| MAC24 | A 24-bit MAC address |
| MAG | Mobile Access Gateway |
| MBR | Maximum Bit Rate |
| MC | Mesh Controller |
| MF | Multi-field |
| MFIB | Multicast Forwarding Information Base |

| | |
|---|---|
| MF-TDMA | Multi-Frequency Time Division Multiple Access |
| MIB | Management Information Base |
| MLD | Multicast Listener Discovery |
| MMT | Multicast Mapping Table 1st generation |
| MMT2 | Multicast Mapping Table 2nd generation |
| MPEG | Moving Pictures Expert Group |
| MSB | Most Significant Bit |
| MSS | Maximum Segment Size |
| MTU | Maximum Transmission Unit |
| NAPT | Network Address Port Translation |
| NBMA | Non-Broadcast Multiple Access |
| NCC | Network Control Centre |
| NCR | Network Clock Reference |
| NE | Network Element |
| NGN | Next Generation Networks |
| NIT | Network Information Table |
| NLID | Network Layer Information Descriptor |
| NM | Network Manager |
| NMC | Network Management Centre |
| NMS | Network Management System |
| OAM | Operations Administration Maintenance |
| OAM&P | Operations Administration Maintenance Provisioning |
| OCS | Online Charging System |
| ODU | OutDoor Unit |
| OFCS | Offline Charging System |
| OID | Object ID |
| ONID | Original Network ID |
| OSPF | Open Shortest Path First |
| OSS | Operations Support System |
| OUI | Organisationally Unique IDentifier |
| OVN | Operator Virtual Network |
| PBA | Proxy Binding Acknowledgement |
| PBU | Proxy Binding Update |
| PCC | Policy and Charging Control |
| PCEF | Policy Control and Charging Enforcement Function |
| PCP | Priority Code Point |
| PCRF | Policy Control and Charging Rules Function |
| PDN | Packet Data Network |
| PDP | Packet Data Protocol |
| PDU | Protocol Data Unit |
| PEP | Performance Enhancing Proxy or Policy Enforcement Point |
| PHB | Per-Hop Behaviour |
| PID | Packet IDentifier |
| PIM | Protocol Independent Multicast |
| PIM-SM | Protocol Independent Multicast - Sparse Mode |
| PLMN | Public Land Mobile Network |
| PM | Performance Management |
| PMIP | Proxy Mobile Internet Protocol |
| PPDU | Payload-adapted Protocol Data Unit |
| QCI | QoS Class Identifier |
| QoS | Quality of Service |
| RA | Random Access |
| RAA | Re-Authentication Answer |
| RA-AC | Random Access Allocation Channel |
| RAR | Re-Authentication Request |
| RBDC | Rate Based Dynamic Capacity |
| RC | Request Class |
| RCS | Return Channel via Satellite |
| RCST | Return Channel via Satellite Terminal |
| RCST2 | Return Channel via Satellite Terminal 2nd generation |
| RF | Radio Frequency |

| | |
|---|---|
| RFC | Request For Comments |
| RIB | Routing Information Base |
| RIP | Routing Information Protocol |
| RL | Return Link |
| RLE | Return Link Encapsulation |
| RMT | RCS Map Table |
| RNC | Radio Network Controller |
| RO | Read Only |
| RPF | Reverse Path Forwarding |
| RPLS | Receiver Physical Layer Segment |
| RRM | Radio Resource Management |
| RTD | Round Trip Delay |
| RTO | Retransmission TimeOut |
| RTT | Round Trip Time |
| RW | Read Write |
| Rx | Receive/Receiver |
| SA | Service Aggregate |
| SACK | Selective ACKnowledgement |
| SCADA | Supervisory Control And Data Acquisition |
| SCPC | Single Channel Per Carrier |
| SCTP | Streaming Control Transport Protocol |
| SDDP | Software and Data Distribution Protocol |
| SDF | Service Data Flow |
| SDP | Session Description Protocol |
| SDU | Service Data Unit |
| SI | Service Information |
| SIP | Session Initiation Protocol |
| SLA | Service Level Agreement |
| SMSS | Sender Maximum Segment Size |
| SNMP | Simple Network Management Protocol |
| SNO | Satellite Network Operator |
| SNR | Signal to Noise Ratio |
| SOHO | Small Office Home Office |
| SP | Service Provider |
| SPR | Subscription Profile Repository |
| SPT | Satellite Position Table |
| SSL | Secure Sockets Layer |
| SSM | Source Specific Multicast |
| ST | Satellite Terminal |
| SVN | Satellite Virtual Network |
| SVN-ID | Satellite Virtual Network IDentifier |
| SVNO | Satellite Virtual Network Operator |
| SW | Software |
| SYN | SYNchronisation |
| Sync | Synchronization |
| TBTP2 | Terminal Burst Time Plan 2nd generation |
| TC | Traffic Class/Classifier |
| TC/PHB | Traffic Class/Per-Hop Behaviour |
| TCP | Transmission Control Protocol |
| TCPM | TCP Maintenance and Minor Extensions |
| TDF | Traffic Detection Function |
| TDM | Time Division Multiplexing |
| TDMA | Time Division Multiple Access |
| TFO | TCP Fast Open |
| TIM | Terminal Information Message |
| TIM-B | Terminal Information Message – Broadcast |
| TIM-U | Terminal Information Message - Unicast |
| TIMU | Terminal Information Message Unicast |
| TIMu/TIM-U | Terminal Information Message – Unicast |
| TLS | Transport Layer Security |
| TMN | Telecommunications Management Network |
| TTL | Time To Live |

| | |
|---|---|
| TV | Television |
| Tx | Transmit/Transmitter |
| TXID | Transmitter IDentifier |
| UDP | User Datagram Protocol |
| UE | User Equipment |
| UMTS | Universal Mobile Telecommunication System |
| USM | User-based Security Model |
| VACM | View-based Access Control Model |
| VBDC | Volume Based Dynamic Capacity |
| VLAN | Virtual Local Area Network |
| VoIP | Voice over Internet Protocol |
| VPN | Virtual Private Network |
| VRF | Virtual Routing Forwarding |
| WAN | Wide Area Network |
| WCDMA | Wideband Code Division Multiple Access |
| WFQ | Weighted Fair Queuing |
| XML | eXtended Marked-up Language |

# 4 Reference system architecture

Refer to [i.1] for the description of a reference system architecture.

# 5 IP routing with OSPF over the satellite interface

IP routing is a key feature to realize seamless integration with terrestrial access networks and to attain multi-vendor interoperability in interactive broadband satellite networks.

This clause provides review and recommendations for IP routing support in satellite networks over the satellite interface with emphasis on the Open Shortest Path First (OSPF) protocol for IPv4 (also called OSPFv2) [i.5]. OSPF performance is analysed under the different modes specified by the OSPF specification, namely, broadcast and non-broadcast modes. Satellite adaptations are proposed to improve OSPF performance in interactive broadband satellite networks. OSPF for IPv6 (also referred to as OSPFv3) [i.7] is also analysed to identify the required adaptations to improve its performance over interactive broadband satellite networks.

The basic IP routing function forwards unicast packets according to the Routing Information Base (RIB). Compiling the information to form a Forwarding Information Base (FIB) can optimize this forwarding. The FIB is usually derived from routing information disseminated in the IP control plane, and stored in the RIB. The RIB is populated either via static configuration or dynamic routing protocols. Dynamic routing protocols are divided into interior (intra-domain) and exterior (inter-domain) protocols. The most common interior routing protocols are the Routing Information Protocol (RIP) [i.6] and Open Shortest Path First (OSPF) [i.5] and [i.7]. The most common exterior routing protocol is the Border Gateway Protocol (BGP) [i.8].

Although all routers provide routing functions, there are significant differences in their feature sets, often determined by where they are placed within the network. For example, customer routers often have reduced feature sets and may use static routing while enterprise routers have expanded feature sets, usually use OSPF and may support VPNs, VLANs, access control, network management, VoIP services, firewall functions and L2 virtualization.

Satellite access networks have typically used static routing, although dynamic routing has been used for backhaul and restoration services in provider networks.

## 5.1 IP routing in satellite networks

A satellite network may support standard IP routing protocols using a L2 interface directly connecting external routers. Alternatively, the routing functions – adapted for satellite networks – may be integrated in the RCST.

Dynamic routing for satellite is defined as the case where a routing protocol is used to route to the networks that are connected via routers to the LAN interface of a RCST. In contrast, static routing uses configuration information in RCST to determine the routing. Dynamic routing is attractive when the routing information is imported from connected networks.

Many network operators implement policies to control the imported routes to protect the routing information in the RIB, although this is not recommended for link-state protocols, such as OSPF. Instead, BGP is widely used for exterior Gateway routing (between domains). BGP use is often accompanied by significant operator policy/configuration, and a need for integration of other protocols (e.g. tunnel management). As such, it may be more appropriate to place BGP functionality in an externally attached dedicated BGP router, rather than within a RCST. Such a router could accept exported routes from the satellite domain.

## 5.2 Packet forwarding in satellite networks

An RCST may forward packets in one of the following modes:

**Static IP routing mode**: The RCST acts as a router within an IP network. Forwarding is performed using the RIB or the compiled FIB. Static IP routing does not require a routing protocol; the routing information is derived from configuration and statically loaded into the RIB. In a star network, the RIB/FIB at a RCST normally has a default route that points to the Gateway, and configured routes at the RCST and GW to the networks connected via the RCST and GW LAN interfaces.

**Dynamic IP routing mode**: The RCST acts as a router within an IP network using a dynamic routing protocol to populate the RIB.

**Dynamic virtual IP routing mode**: A router that supports VRF groups maintains a set of completely isolated routing entities, one for each supported VRF group.

Dynamic routing protocols, such as OSPF use link-local multicast, however many currently deployed satellite networks do not support this. Although the satellite outbound may allow RCSTs to receive multicast, the Gateway often does not replicate inbound multicast to all RCSTs. Internet links that do not support multicast routing should use point-to-point mode between adjacent routers. This mode cannot take advantage of the improved transmission efficiency offered by multicast.

Static IP routing is the standard mode recommended for stub networks: static routing information may be distributed in the satellite network by configuration. It may also be possible to export routes from the Gateway to the RCSTs using standard OSPF routing packets.

Support for dynamic routing requires correct treatment of link-local multicast (and at minimum, support for reception of these packets from the RCST by the Gateway/NCC).

## 5.3 Satellite network routing topologies

Figure 5.1 shows two example network topologies for OSPF.

- Stub Networks. In this topology, Figure 5.1a, the RCSTs are stub routers that correspond to the scenarios normally expected for Consumer, SOHO, Multi-dwelling and Backhauling profiles. This topology assumes that the RCST is aware of (or allocates) the address space for the network connected via the RCST LAN interface. In this case, the NCC/Gateway already knows the remote network topology. This is expected to be often the case for star networks using private addressing, or where the LAN interface connects a remote NAPT or provides a point-to-point link that carries a tunnel (as in many Backhauling scenarios).

- Routed Stub-Networks. In this topology (Figure 5.1b), an RCST may have independently addressed sub-networks connected to the RCST LAN interface. This topology is expected to be common for Government and Corporate/Institutional use. It assumes that the RCST connects independently managed networks, where the address space for the network connected via the RCST LAN interface is not under the control of the NCC. In this case, the remote network topology is not directly known by the NCC and routing information is required to indicate which network address is reachable via which RCST. Dynamic or static routing is used to inject the routing information into the RIB at the satellite NCC/Gateway. The remotely connected networks may set a default route delivering all non-local traffic over their RCST air interface.

- Routed Dynamic Networks. This topology consists of sub-networks connected to the LAN interface, and is the same as previous with one exception (Figure 5.1b). It connects networks using a dynamic routing protocol within the attached network. This is needed for networks that employ dynamic routing (e.g. to realize alternate paths to the satellite network, or where the satellite/alternate path is used as a backup for restoration of service following a failure). The important difference here is that the RCST should fully participate in the routing protocol exchanges (i.e. routing updates export topology information about the satellite network and import information about the reachable networks via the LAN interface). The key difference of this topology is that policies are needed for route import/export and whether the RCST router functions as a routing border.

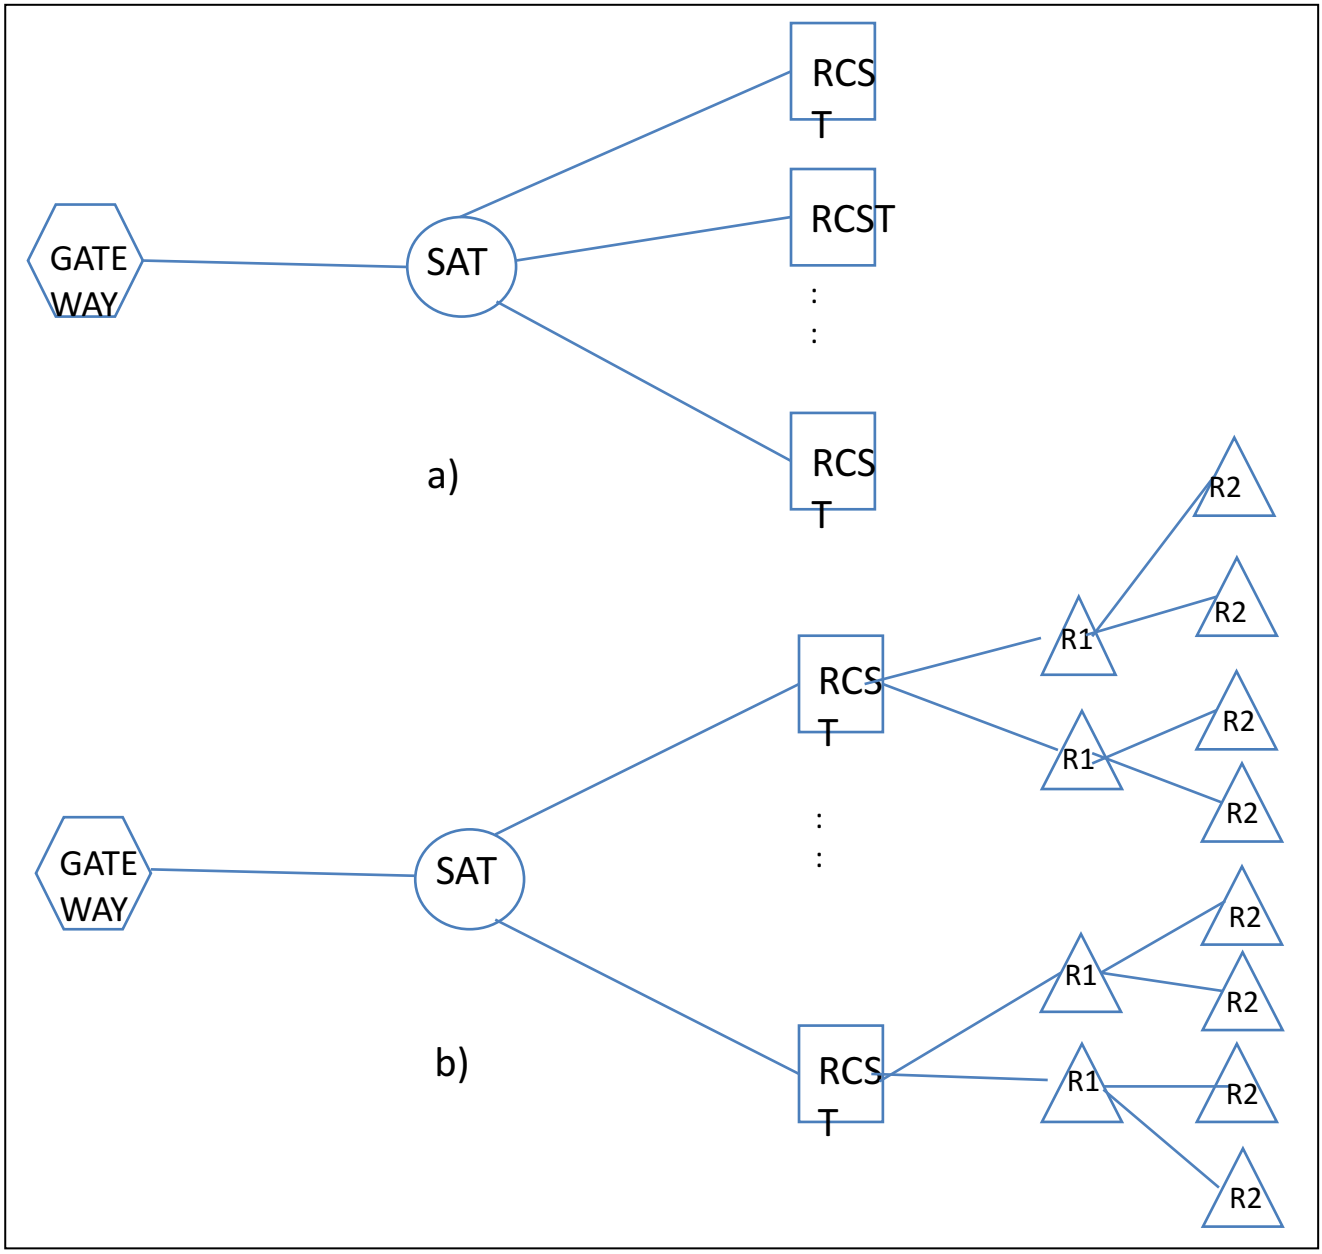


**Figure 5.1: Example network topologies**

The use of dynamic routing in stub routers for a satellite star network is not usually necessary, since each RCST may be preconfigured with a static default route to the Gateway and routes to each locally attached network. The Gateway/NCC RIB would be statically configured to support the address range delegated to each RCST.

Dynamic routing is desirable in cases where the set of networks reachable via an RCST can change without reconfiguration of the NCC. This may be the case for example when alternate paths exist (e.g. the connected network is also reachable via a terrestrial (backup) link) or where the NCC does not control the addressing plan (e.g. the operator of the remote network can move assigned IP addresses between sites without reconfiguration of the satellite network).

## 5.4 Dynamic routing using OSPF in transparent star networks

This clause describes the use of OSPF for dynamic IP routing in a satellite star network and provides guidelines on its use.

The core OSPF algorithms are election of the Designated Router (DR), flooding of routing information, and OSPF route calculation. OSPF packets are directly encapsulated in the IP protocol using protocol number 89, using a combination of unicast and link-local multicast.

OSPF sends Hello packets periodically on each interface to discover, establish and maintain a router's neighbour relationships. This also facilitates OSPF router configuration by indicating any support for optional capabilities. Routers of different capabilities can be mixed within an OSPF routing domain using the features advertised in this field. Note that Hello message exchanges serve to verify router configuration (e.g. adjacency, addresses used and protocol options) in addition to discovering neighbours and electing a DR. This functionality of Hello message exchanges makes it unattractive totally suppressing their transmission in a satellite network.

OSPF supports various modes of operation. It runs in point-to-point, broadcast and nonbroadcast modes (in addition, virtual links can also be configured). Over non-broadcast networks, it can operate in one of two modes: non-broadcast multi-access (NBMA) and point-to-multipoint.

Figure 5.2 shows the OSPF architecture in star and mesh systems. The first diagram corresponds to a star network, in which the DR is located in the Hub/GW. The OSPF function should give support for all the existing SVNs. The second diagram illustrates the mesh case, where each SVN has an RCST acting as GW, holding the DR.

More complex architectures may include, in mesh networks, a GW supporting several traffic SVNs, giving OSPF support to these SVNs, resembling the transparent case.

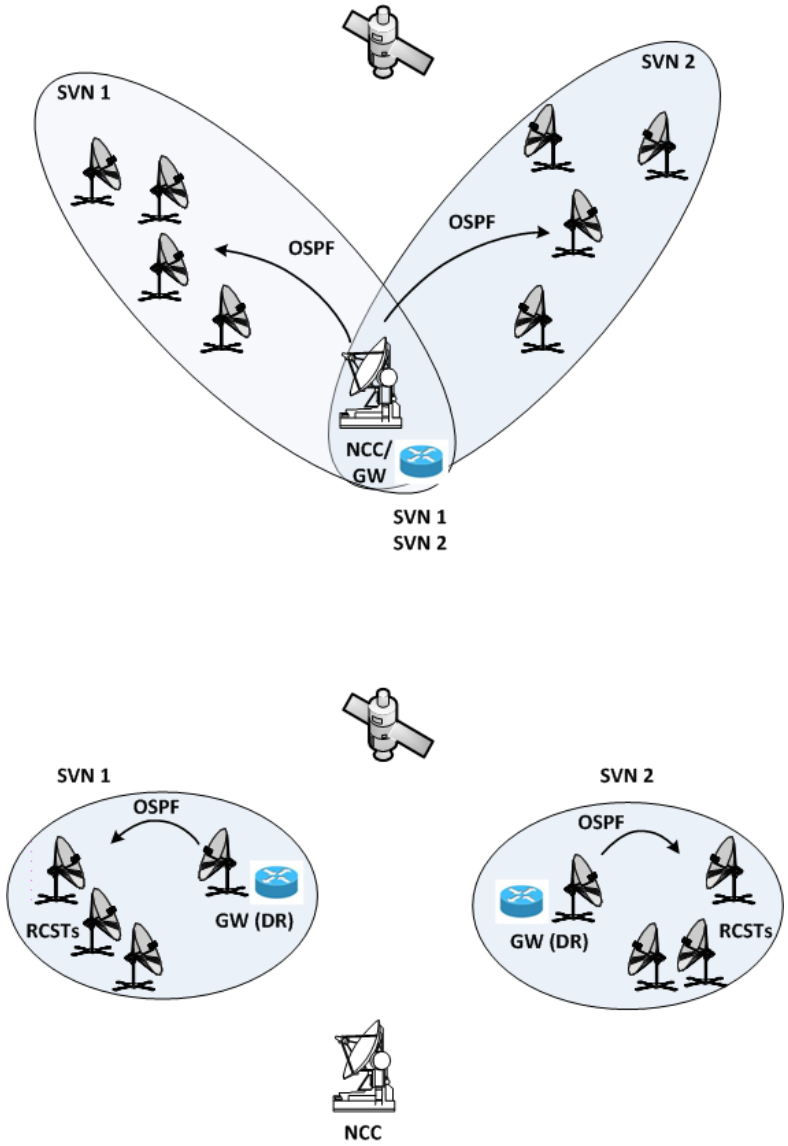


**Figure 5.2: OSPF architecture for star and mesh satellite systems**

## 5.4.1 OSPF for IPv4

OSPFv2 is standardized as [i.5] and its extensions. OSPFv2 uses unicast and link-local IPv4 multicast.

## 5.4.2 OSPF for IPv6

Dynamic Routing for IPv6 is supported using OSPFv3 [i.7]. This standard replaces its earlier version, and removes support for dynamic multicast routing, instead preferring PIM-SM. OSPFv3 relies on IPv6 support on the router interfaces, including support for IPv6 link-local multicast.

OSPFv3 presents some changes compared to OSPFv2 [i.5], mainly to support the increased address size and routing prefixes, and to provide a larger Options field. The core OSPF algorithms (e.g. DR election, flooding, OSPF calculation) remain unchanged. OSPFv3 also supports multiple routing instances on a link. The changes do not significantly impact the performance over satellite networks, the more compact format used in OSPFv3 results in approximately the same overhead even when using the larger IPv6 addresses.

### 5.4.3 OSPF Designated Router

In OSPF broadcast mode, a Designated Router (DR) (and backup DR, BDR) is elected via the exchange of Hello packets. Once elected, the DR sends multicast Hello packets every *HelloInterval* to other routers in the same area, whereas the rest of the routers send unicast Hello packets only to the DR and BDR. In other words, DR and BDR should be adjacent to the rest of the routers of the area. If a router does not receive a Hello packet from a certain adjacency during a *RouterDeadInterval*, it will declare the router down. In the topologies shown in Figure 5.1, in a satellite star network, the DR functions will normally be located at the Gateway/NCC. The BDR role may not be needed (because a Gateway failure results in the network becoming unavailable), or the BDR role may be co-located at another router also at the Gateway/NCC.

### 5.4.4 OSPF – NBMA mode

OSPF supports non-broadcast multiple-access networks (NBMA). In NBMA networks, the DR and BDR are statically configured, that is dynamic discovery of neighbours is not performed in NBMA mode. In a satellite environment, the Gateway would usually be configured as the DR. It should be emphasized that, in NBMA mode, OSPF messages are always unicast; this would include OSPF messages on the forward link, which is an inefficient use of forward link capacity.

### 5.4.5 OSPF – Asymmetric multicast support

A RCST may support link-local multicast transmission towards the Gateway, in which it can send IP multicast packets to the Gateway via the return link. A RCST may send to the multicast groups *AllSPFRouters* (224.0.0.5 for IPv4 and FF02::5 for IPv6) and *AllDRouters* (224.0.0.6 for IPv4 and FF02::6 for IPv6) but, in many current star satellite systems, this does not result in the retransmission on the forward link of these multicast IP datagrams that have originated from a RCST. Therefore, even though each RCST may send an OSPF Hello packet to the IPv4 group 224.0.0.5 (FF02::5 for IPv6), other RCSTs may not receive these packets. This means that RCSTs may be unaware of other RCSTs in the same OSPF area.

In summary, it is recommended that each RCST have the capability to transmit link-local multicast packets to the Gateway, even if the RCST is not updated to support other multicast functions.

### 5.4.6 OSPF RCST steps

Figure 5.3 shows the steps followed by an RCST that implements dynamic routing. The initial OSPF configuration (DR address) is obtained in the logon response message. The Hello protocol allows the creation of adjacencies and OSPF options configuration. During the OSPF flooding process, the RCST synchronizes with the SVN routing information and updates its RIB. Upon a LAN update (new public prefixes reachable through this RCST), the RCST propagates the new routing information towards the DR.

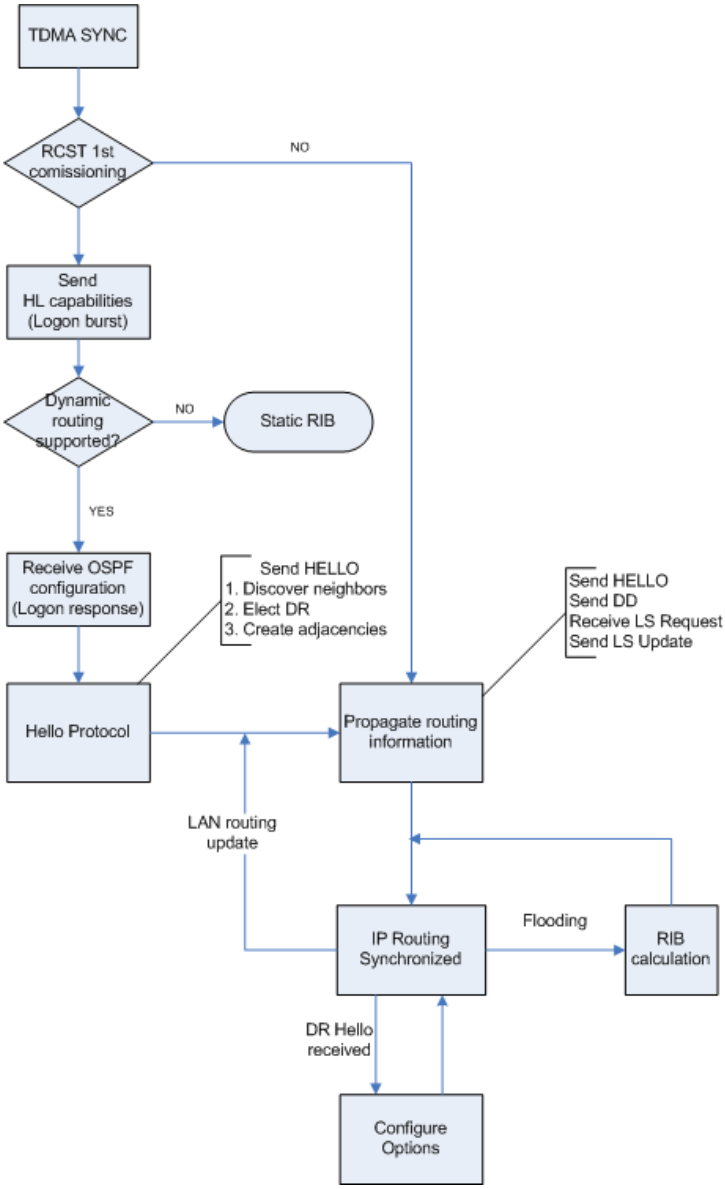


**Figure 5.3: RCST OSPF steps**

## 5.4.7 Optimization of adjacency detection for satellites

In many deployed networks, the default OSPF configuration does not offer sufficiently fast detection of loss of connectivity to an adjacent OSPF router. One method to increase detection of this failure would be to send Hello packets more frequently and reduce the corresponding timers. However, this increases the overhead, and is undesirable for wireless/satellite links.

Another method to quickly detect loss of connectivity to an adjacent router is to run a lightweight UDP protocol, known as Bi-Directional Forwarding Detection (BFD), specified in [i.9] and [i.10]. BFD can enhance detection of failures of an adjacency by providing a signal to the routing engine following a loss of a link. The BFD exchanges, while small, may also be undesirable for a satellite system, since these (like Hello packets) are sent irrespective of the traffic on a return link. In a satellite context, a similar gain to the use of BFD may be achieved through the use of a lower layer signalling mechanism that detects loss of the channel. This solution may save overhead by avoiding frequent packet exchanges.

One optimization could be to eliminate use of Hello packets to increase performance. Loss of router adjancency over a satellite link could be quickly detected at the LL layer (below the HLS) and indicated to the RCST router without the need of IP-level Hello exchanges. Furthermore, a periodic exchange of Hello packets would also consume capacity on the return link, even for a RCST that carries little or no other return link traffic. However, it should be noted that Hello packets are not only used to verify adjacency (lack of which can be detected below IP), but also to notify the configurations supported through their Options field.

The recommended solution is to reduce the periodicity of broadcast Hello packets from the Gateway, and to suppress Hello packets from RCSTs; except, when an RCST starts or restarts the OSPF routing process. This recommendation is motivated by a desire to reduce the capacity consumed by dynamic routing traffic. The cost of transmission on the forward link is much lower than for transmission on the return link from an RCST, which (in any case) could be idle. In addition, a single copy of the Hello packet is multicast to all RCSTs within a satellite virtual network. The periodicity may be reduced, since RCSTs do not need to use this to elect a DR (this is statically configured to be the Gateway), and an RCST may cache the options and the DR address from previous messages. The advantage of this approach is that it preserves the normal characteristics of the OSPF protocol – that is it provides a mechanism at the IP level to detect and confirm the adjacency, options in use, etc. Such confirmation provides logging information to an operator that ensures that incremental deployment of updates and the validity of any changes in configuration are noted at the IP network level (rather than being solely reliant on correct configuration of lower layers).

The router memory requirement seems to be acceptable for the current designs of RCSTs. The satellite capacity consumed will depend on the topology and OSPF mode.

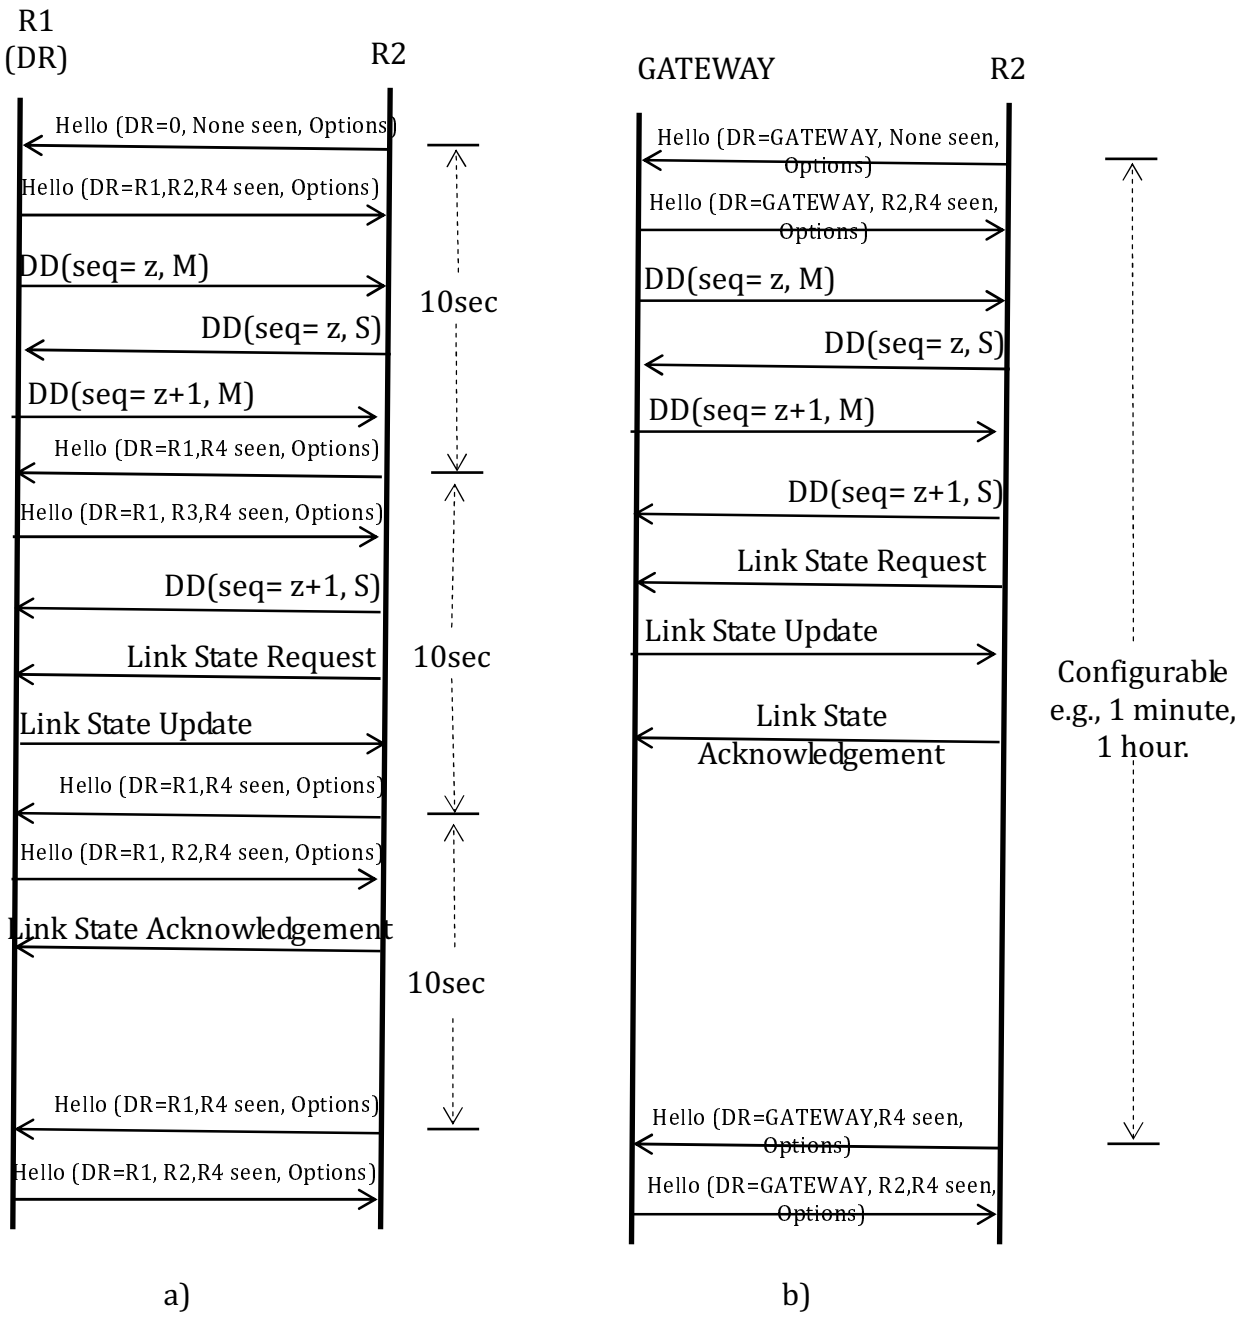


**Figure 5.4: OSPF packet exchanges at startup**

Figure 5.4a) shows the default OSPF packet exchanges between a DR and a router, R2, in broadcast mode. Figure 5.4b) presents the packet exchanges for the satellite-adapted version. The traffic Gateway operates as the DR, and the Hello packets interval is updated to advertise the configuration and options supported. The election of DR in (b) will be suppressed since the Gateway Router will take this role.

## 5.5 Dynamic routing for mesh satellite networking

In a mesh network, an RCST may be capable of direct communication with another RCST, without requiring the use of a Gateway to regenerate the signal. Although mesh communication enables traffic to be directly sent from one RCST to another, the underlying communication may be still under the control of a single NCC.

Mesh networks may be divided into several categories:

- Large Enterprise networks, e.g. with 200-2000 RCSTs controlled by a NCC. Such networks could be used for LAN interconnection, with the ability for any RCST to directly reach any peer RCST. The destination RCST could be identified by a layer 2 label. If this network operates at layer 2, then a next-hop resolution method is required to determine the value of the label to be used to reach a destination RCST. Routing is required to identify which RCST is to be used to reach a destination address assigned to a remote network.

- Hybrid star/mesh, where traffic is routed via either star or mesh connectivity. An RCST could direct traffic via either a traffic Gateway or a mesh connection, depending on the intended destination. QoS-based routing is also possible. Example applications of using QoS to make routing decisions include routing VoIP over mesh (to minimize delay), data over a star connection (to minimize cost). Routing is required to identify which RCST is to be used to reach a destination address assigned to a remote network, the DR is assumed to be at the NCC/Gateway.

- Small hub-less mesh. In this network, one RCST may assume the role of an OSPF DR (as in a star network, there may be a natural choice of RCST, if one RCST acts as a trafficGateway for the user traffic). This style of network may be well suited to applications such as SCADA networks. From a network perspective, this topology resembles that of a regenerative satellite network, where connectivity between RCSTs is possible without the role of a dedicated Gateway terminal.

- Mesh networks over semi-transparent satellite where a satellite supports multiple types of transponder: one optimized for star networks and one optimized for mesh use. From a network-layer perspective, this resembles the hybrid star/mesh case.

In summary, a dynamic routing mechanism is needed for mesh networks to direct traffic over the mesh connection, since the routes available depend upon the mesh capabilities. This mechanism is also needed to integrate mesh networks with terrestrial networks using OSPFv3.

A possible solution may be to use standard IP-based methods for routing (as described in the present document) and to synchronize the detection of router adjacency with the lower layer functions that establish and release a physical layer stream.

Two design options are presented to realize such a system:

1) The NCC is responsible for setup and clear down of mesh connectivity, and one approach could be for the co-located DR router to make all mesh routing decisions following establishment of a lower-layer mesh connectivity by the NCC. The establishment of a mesh connection would therefore trigger an OSPF routing update.

2) Another option is to establish the mesh connectivity at the physical and link layers via the NCC, and then for the RCST to exchange Hello packets over the established link, resulting a routing update from the RCSTs to the DR, and a corresponding routing update to all RCSTs. This interaction follows normal OSPF behaviour. The exchange of the Hello packets also establishes that the link is operational and permits the exchange of configuration data via the Hello packet options. Reachability can be validated by cross-layer mechanisms, eliminating the need for periodic Hello packets over an established mesh connection.

Method-1 eliminates some RCST signalling, although it is less robust than method-2, since it does not validate the IP path. Method-1 also requires that any RCST policy (e.g. which addresses/traffic classes are to be routed via the mesh connection) needs to be configured and maintained at the NCC, rather than allowing the possibility that this could be a locally-configured RCST policy. Dynamic Connectivity Protocol (DCP) explained in [i.1] is recommended for setup and release of lower-layer mesh connections.

Figure 5.5 shows an SVN routing update triggered by a routing change in the LAN connected to an RCST. This scenario requires multicast transmission of LSA updates from at least the RCST with the role of DR in the SVN. An RCST logon event should result in sending a Hello message and may also trigger the sending of routing updates by the DR.

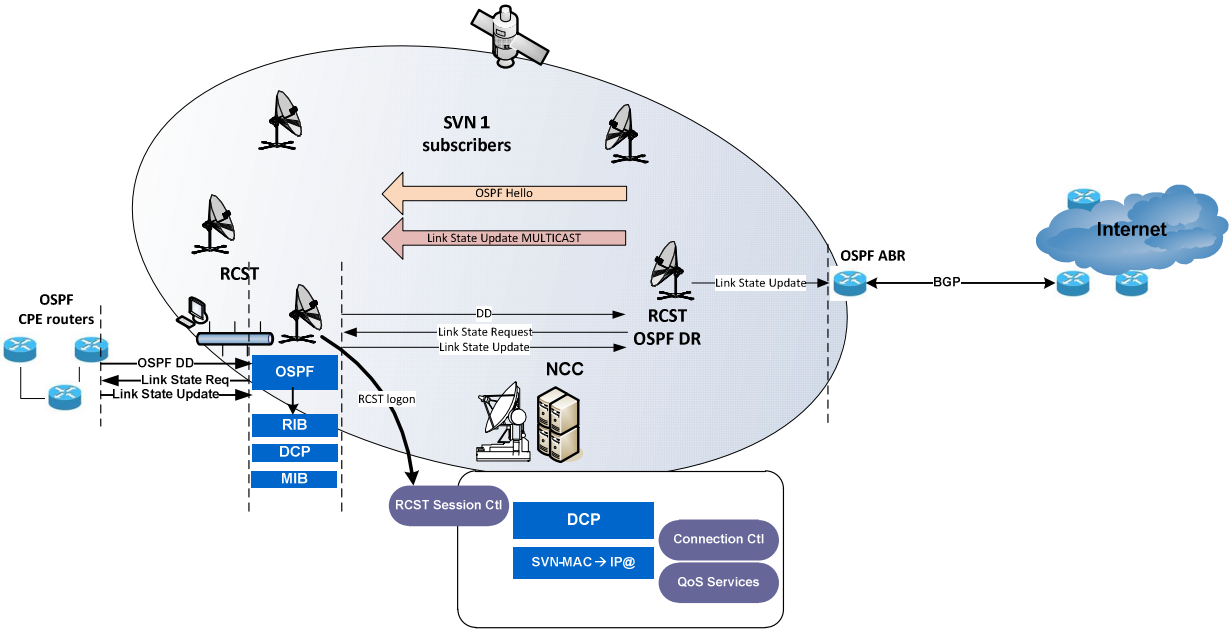


**Figure 5.5: OSPF in mesh network**

The corporate scenario is characterized by many terminals and a medium-to-big GW including the DR. The RCST-GW LAN interface address is configured as the default IP next hop in all RCSTs of the SVN. This scenario may support multiple SVN groups by the GW. In case a dedicated RCST-GW is not used per SVN, the common RCST-GW should support OSPF independently on each VRF group.

When the processing capabilities of the RCST-GW are not enough to host the OSPF router, OSPF forwarding capability by the RCST-GW towards the DR may be requested, unless the GW includes several RCST-GWs and the satellite AS is divided into smaller areas. This means forwarding OSPF packets to/from the satellite interface without decrementing the TTL field in the IPv4 header.

In this scenario, the GW OSPF router is integrated with terrestrial networks, being the satellite network Area Border Router (ABR) and using BGP protocol.

In a mesh network, it may be possible having a secondary (backup) DR, corresponding to a second RCST-GW, in an SVN. The OSPF routers may be or not co-located. In the example of having a GW including two RCST-GWs in active/stand-by configuration, when the RCST-GW1 fails, OSPF should restore mesh links between the RCSTs in the SVN and the RCST-GW2, which is the backup DR.

## 5.5.1 OSPF and DCP

In mesh systems, OSPF packets are propagated over the mesh links established via DCP. For systems where the OSPF function is activated on the satellite interface, DCP only provides L2 address resolution function, and the IP routing information needed to construct the request messages should be provided by the source RCST, using the information in its RIB (dynamically updated by OSPF). A new DCP request will be issued by the RCST when there is no other mesh link opened directed to the same destination IP, in the same SVN, and using the same HL service. In other case, IP packets are forwarded to one of the opened links.

Figure 5.6 shows a successful DCP exchange. An IP packet reaches the LAN interface of RCST1. Thanks to its RIB, RCST1 knows that the next hop IP address to reach the packet destination is the RCST IP router address and includes it in the DCP request message.

If the source RCST supports multiple SVNs, the SVN mask included in the request should be used by the NCC to guess infer the SVN in which the mesh link is to be established. With this information and the destination IP address (or the next hop IP address), the NCC is able to locate the MAC24 address of the peer RCST, which will be included in the DCP response message.

If the next hop IP address of an outgoing packet is not found in the AR database, a DCP Link establishment request is triggered by the RCST to find the L2 address of the next hop. In case that the system does not support the dynamic routing function (e.g. OSPF), the DCP protocol can assist the RCST with IP routing information.

The NCC allows establishment of DCP Links only between RCSTs belonging to the same SVN or located in a common VRF domain, otherwise rejecting the Link requests.

The RCST may indicate in the request message the next hop IP address (Next hop address field in the Triggering datagram identifier IE) according to its RIB. When this field has been filled by the RCST and the NCC cannot identify the destination RCST from the triggering packet destination address, then the NCC should use the address of the next hop field to obtain the MAC24 and the FPDU identifiers corresponding to the peer RCST.

The transparent mesh RCST obtains the bursts to be decoded from its Active Links Table. The information about the other RCST Assignment_ID is obtained from the DCP messages sent by the NCC. It is assumed that, in a mono-beam transparent mesh system, all the RCSTs decode the same TBTP2 and therefore can extract information about the timeslots used by the other peer of the mesh link.

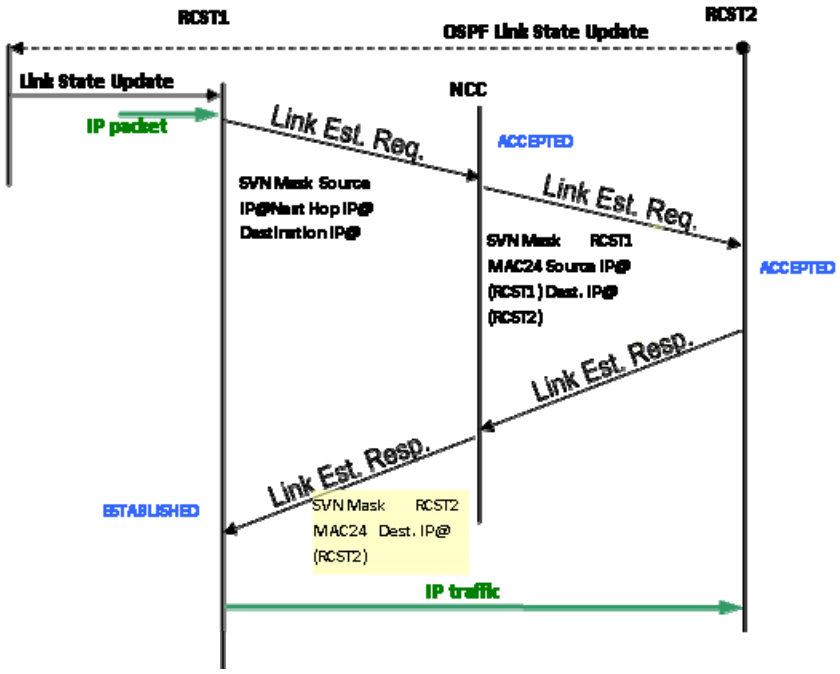


**Figure 5.6: Example of DCP exchange for bidirectional mesh link (addressing parameters)**

Figure 5.7 shows the internal procedures of an RCST supporting dynamic connectivity. The RCST should issue DCP requests when the next IP hop (corresponding to IP packet destination) is not found in the DCP active Links table. When found, the associated SVN number should match as well.

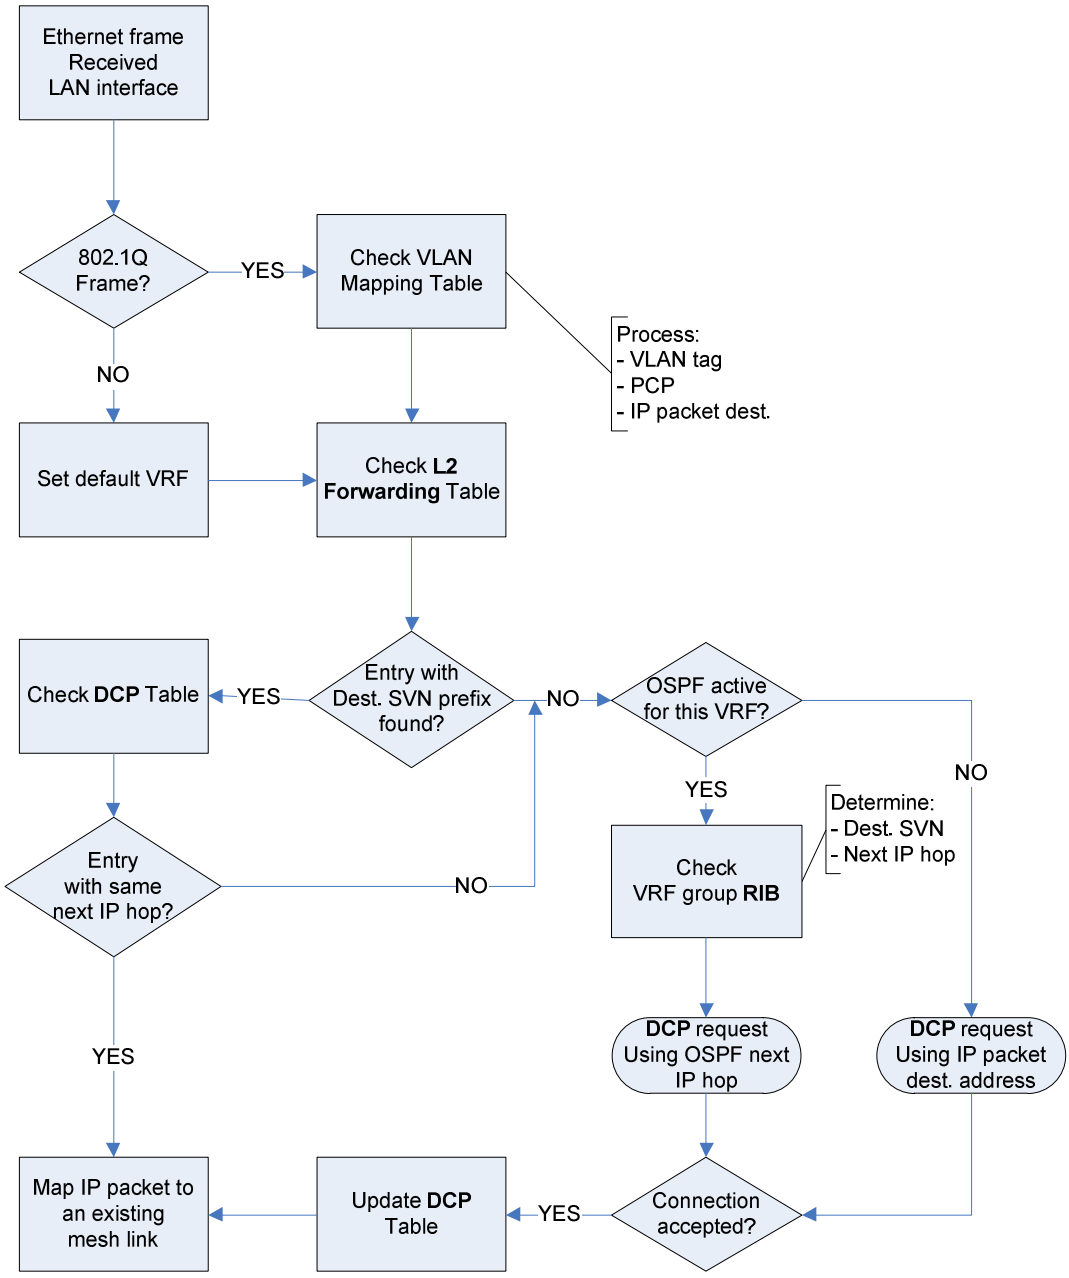


**Figure 5.7: IP routing and DCP**

## 5.6 Recommendations for satellite routing support

This clause describes dynamic IP routing using OSPF in a satellite star network.

OSPFv2 and OSPFv3 are used for IPv4 and IPv6-based networks, respectively. An implementation should adopt the standard mechanisms specified by the IETF, together with the updates recommended in this clause.

Two topologies can be considered for transparent star satellite networks where the Gateway takes the role of the DR:

1) Use of OSPF may not be desirable for large satellite networks where all the RCSTs form one routing area. The advantages of dynamic routing are minimal where the addressing plan at the sites connected via the RCST LAN interface are expected to be under the control of the NCC. This is normally expected to be the case for Consumer, SOHO, Multi-dwelling and Backhauling scenarios. Therefore, static routing may be preferred in these cases. An RCST may be preconfigured with a static default route to the Gateway and routes to its locally attached networks. The routing tables at the traffic Gateway could also be statically configured to support the address range delegated to each RCST. Configuration of static routes could be made through standard IP management (e.g. SNMP, netconf, CLI), or could be considered as a part of the lower layer configuration. In IP networks, the former is usually preferred.

2) Dynamic routing using OSPF may be desirable for scenarios where the satellite network feeds one or more networks (e.g. using public address space) where the addressing plan is not under the control of the NCC. It is also desirable for networks that employ dynamic routing (e.g. to realize alternate paths to the satellite network, or where the satellite/alternate path is used as a backup for restoration of service following a failure). This is expected to be common for Government and Corporate/Institutional use.

It is recommended that a star satellite network uses the OSPF Broadcast mode to take advantage of the lower transmission cost of sending multicast packets from the DR during the flooding process. This requires that link-local multicast packets originating from an RCST are sent on the inbound (return) link to the Gateway/NCC where the OSPF process executes. The remainder of this clause makes recommendations for optimizing performance in a satellite network.

### 5.6.1 Recommendations for transmission of Hello packets

This clause updates the Hello processing described in [i.5] and [i.7]. This update applies only to a satellite interface. The reception of Hello packets by an RCST should be used to verify the correctness of the present configuration. It is desirable that IP functions are verified at the IP layer, rather than being entirely reliant on the correct configuration of end-points and lower layers. This promotes the Internet concept of "fate sharing", whereby a network path can be used if it is known to be functional, rather than relying on correct pre-configuration of protocols at lower layers.

Hello packets should be sent by the Gateway using multicast to reach all RCSTs within a satellite virtual network. The periodicity of these messages may be reduced depending on operator needs, but should not be reduced to less often than every 30 minutes (a value chosen as a trade-off between ability to detect misconfiguration and overhead).

Hello packets received by an RCST should be used to verify the correctness of the present OSPF routing configuration, including checking the DR IP address and options. It is desirable that IP configuration values are verified at the IP layer, rather than being reliant on correct configuration of end-points and lower layers.

An RCST should cache the last received Hello packet. This optimization may enable it to quickly restart the OSPF process with a relevant configuration. If used, this cache should be cleared when the NCC restarts the satellite network following a configuration change that affects IP routing (e.g. change of addressing plan or change in Gateway router configuration).

It is recommended that an RCST sends a single Hello packet to the DR during the "set-up" phase of an RCST. This packet indicates the router's capabilities through the Options field. Hello packets may also be exchanged at periodic intervals to verify the state has not been changed, but this interval need not be small (as in reachability detection), and a much longer value may be configured (e.g. at intervals of minutes/hours). Note that sending this packet once does not constitute a major overhead, but confirms reachability. A drawback is that this packet could be lost in the network i.e. the configuration soft-state is not refreshed by the protocol. Network operators need to be aware of this possibility when managing their networks. This does not impact dynamic routing, provided that the configuration of RCSTs and Gateway are consistent. The RCST should resend this packet if the IP routing configuration of an RCST changes or the OSPF process is restarted.

An RCST that detects a loss of the forward link or a state transition at the lower layer that prevents IP transmission on the return link should update its OSPF adjacency as if the router at the RCST had failed to receive a Hello packet. This failure detection is similar to the use of BFD in [i.71].

### 5.6.2 Recommendations for routing topology update

It is recommended that standard OSPF methods are used to propagate the routing information, and RCSTs are enabled to send these updates using link-local IP multicast to the Gateway router.

### 5.6.3 Recommendations for defining OSPF Areas

The signalling cost of using dynamic routing is highly dependent on the topology of the network. Scaling, i.e. the number of routers using a single DR needs to be considered to minimize routing traffic. Judicious configuration of border routers (ABRs) to divide routing areas is recommended at an RCST, where the RCST connects more than a few OSPF routers via the LAN interface. This may be especially useful when the routing information may be summarized or when there are frequent routing updates within the network fed by an RCST router LAN interface. The principles for configuring ABRs are not specific to satellite and advice can be obtained from usage in other networks.

# 6 IP multicasting

This clause elaborates on the MMT2 supported method for mapping multicast to L2 as being the most versatile method of the two alternatives specified by [i.3].

## 6.1 Mapping IP multicast to L2

### 6.1.1 Multicast over Ethernet

Modern Ethernet controllers filter multicast frames out of received frames to reduce the host CPU load. This is achieved by deriving a L2 multicast group destination address for each IP multicast group that needs to be received/forwarded. The set of active L2 addresses is stored in a table maintained by the host software. This table is used to decide whether a multicast frame received on an interface is forwarded to L3 or discarded. The L3 address is also checked against the set of groups to be received before the packet is forwarded to higher layers for further processing.

### 6.1.2 Mapping of IPv4 addresses

The mapping between IP and L2 multicast group addresses at the LAN interface is usually provided by a standard method [i.11] that derives a L2 address for each IPv4 multicast group destination address. This method is commonly used for all multicast networks, and it is also the method specified for GSE satellite systems that use the 6-byte address format. [i.3] specifies two alternatives for mapping multicast to MAC24, one operating without support of MMT2 and one using MMT2. The NCC determines which method that applies, and provides the necessary configuration as part of the LL service configuration.

### 6.1.3 Mappings for IPv6 address

IPv6 includes multicast as a standard function, and many core IPv6 protocols rely upon multicast support. The same IPv4 multicast requirements exist for IPv6. However, there are some additional considerations:

1) IPv6 multicast redefines the way scoped addresses are specified, this places additional constraints on filtering addresses when forwarding between different networks (or virtual networks).

2) Some protocols (e.g. neighbour discovery) generate link-local IPv6 multicast addresses, which means that many LANs carry a large range of IPv6 multicast groups, none of which is intended to be forwarded by a router.

3) The IPv6 address range is larger than that of IPv4, which can result in more overlap of addresses (i.e. two addresses in different address blocks can map to the same L2 address).

For dual-stack deployments with significant levels of both multicast IPv4 and IPv6 traffic, it is recommended that separate L2 address spaces are used for the two services to avoid address overlap. In cases with lower levels of multicast traffic, or where the addressing plan is under the control of the operator (who could assign addresses to avoid overlap), the two protocols (IPv6 and IPv4) may share the same L2 address space.

[i.3] specifies two alternatives for mapping multicast to MAC24, one operating without support of MMT2 and another using MMT2. The NCC determines which method that applies, and provides the necessary configuration as part of the LL service configuration.

## 6.2 Operator-controlled mapping of Layer-2 multicast group addresses

The mapping between L3 and L2 addresses specified in [i.11] may result in two L3 groups may map to the same L2 address, which is usually referred to "address overlap". This is not normally a concern for LANs thanks to the additional IP multicast group address filtering at the IP layer. It can be a significant issue when two completely different services map to the same address, since the forwarding is usually controlled per L2 address. This can, for instance, occur when the network link carries traffic from more than one service operator (e.g. in a multicast Internet exchange point).

One way to avoid the issue of address overlap is by careful L3 address assignment. This is recommended in static configurations although hard to manage with dynamic multicast. It is common for applications to choose well-known IP multicast addresses, and this would result in unexpected behaviour when more than one virtual addressing space is used. Separation of different multicast address scopings is essential for proper multicast operation [i.12].

DVB-RCS2 may provide support for multi-operator use of multicast. This addresses a need to support Internet multicast access or/and when multiple virtual networks are supported over the satellite, by allowing a satellite virtual network operator to control the mapping to L2 address. The mappings are configured in the Feeder (along with any required QoS requirements). This device is managed by an SNO and also coordinates the mapping for unicast network-layer packets for each SVNO. The mappings configured at the Feeder are also announced by using a control table called the DVB Multicast-Mapping Table (MMT2). This is organized into a set of sections directed to each SVN that supports multicast.

## 6.3 IP multicast forwarding over satellite and LAN interfaces

This clause explains the use of IP multicast control techniques and the use of the Internet Group Management Protocol (IGMP) [i.13], [i.14] and [i.15] to deliver multicast content to an RCST. Equivalent behaviour is also expected for IPv6 using the Multicast Listener Discovery Protocol (MLD) [i.15], [i.16] running over ICMPv6. Support for IGMP over the LAN interface can be classified as either passive or active. For DVB-RCS2, the active mode refers to use of an IGMP or MLD intercepting proxy agent operating over the LAN interface at the RCST.

Multicast receiver hosts do not participate in routing decisions, and instead use multicast control protocols to signal the set of groups that they wish to receive. IGMP is used for IPv4 and MLD is used for IPv6. Upon reception of these control messages, a multicast router triggers appropriate routing messages (e.g. PIM-SM messages [i.17]) to control forwarding from any routed upstream network node that supports IP multicast.

Managed Ethernet switches typically implement an IGMP Proxy [i.18], in which an agent intercepts membership reports from hosts and uses this information to determine over which LAN interfaces to forward IP multicast packets. This is similar to the way a multicast-enabled IP router processes these reports.

An RCST may be configured to function in either a passive or an active multicast control mode in regards to how they forward multicast traffic on its LAN interface. In passive mode, no IGMP messages are processed by the RCST and no multicast membership reports are sent on the satellite interface, whereas in active mode IGMP messages are terminated at an IGMP proxy agent in the RCST and may then be sent over the satellite (when dynamic multicast forwarding is used). In a satellite network, an RCST implementing active mode over the LAN interface also uses a proxy agent that participates in the multicast control protocols, to populate a local data structure identifying the set of presently active IP multicast groups (the Multicast FIB, MFIB). In contrast, a RCST using passive mode over the LAN interface provides a configuration interface to insert entries in the MFIB. When using passive mode, an RCST forwards traffic to the LAN interface independent of whether there is an active receiver on a connected host. This mode resembles broadcast, in that the Service Provider determines the forwarding. In many cases it is not necessary that all multicast groups are forwarded to the LAN interface by all RCSTs, and in passive mode this is controlled by RCST configuration. Downstream devices connected to the LAN Interface of an RCST such as managed Ethernet switches, may control the propagation of specific groups using standard methods such IGMP snooping [i.19] or IGMP proxy [i.18]. IGMP Proxy is generally preferred.

The procedures for multicast traffic sources attached to RCSTs and bidirectional multicast are not defined in DVB-RCS2. The current document also does not specify multicast router interactions with an RCST. Neither does it describe the support of multicast routing over the satellite interface.

The preceding text focuses on multicast forward control at the RCST over the LAN interface. The remaining text delves into the forwarding multicast traffic that is received on the satellite interface.

The NCC is responsible for all multicast transmission on the forward link. Prior to multicast data transmission on the forward link, the NCC should first configure the Feeder and the Gateway Router with entries for those multicast streams that are to be forwarded. Once configured, the Gateway Router joins the requested group at its upstream interface. The feeder encapsulates the multicast traffic and forwards it over the satellite air interface. Static and dynamic forwarding are distinguished in this case: In static forwarding, this process is completed before the RCST needs to receive the multicast group. In dynamic forwarding, this process is completed the first time an RCST requests to join a specific group. An RCST that wishes to receive multicast traffic with a specified IP address on the forward link should first construct a layer 2 filter containing the GSE labels with which the multicast traffic is sent. This table may be directly mapped using the information in the MMT2 to identify the GSE address used to carry a multicast flow. Once configured, the filters forward all traffic with the label to the IP layer where the RCST filters the traffic based on the IP network layer address using the information in the MFIB.

Static forwarding on the satellite interface may be extended by enabling a proxy agent at the RCST to provide active mode at the LAN interface. Static forwarding on the satellite interface and active mode on the LAN interface is intended to be the default for DVB-RCS2 RCSTs.

The proxy agent intercepts a group management protocol (e.g. IGMP, MLD), by intercepting packets received on the LAN interface to build a local forwarding table (e.g. held in the Multicast Forwarding Information Base, MFIB). Multicast traffic received from the forward link is only forwarded to the LAN interface when there are active receivers connected via the LAN interface that need to receive the specific IP group. This prevents the LAN from having to carry traffic for services for which there are no receivers, an attractive optimization provided in most multicast networks.

Static forwarding over the satellite has advantages in terms of simplicity of design of the Gateway and control of the QoS offered to each multicast group. However, the approach relies on the operator determining what content is to be received at any time. While this is appropriate for pre-scheduled transmissions (such as file updates, IPTV broadcast, etc.), it is not appropriate for applications that are user-driven (such as video-on-demand, multi-party conference, collaborative working applications and service discovery). User-driven applications often cannot predetermine the set of multicast groups that will be used, and it is often not feasible to forward all multicast traffic over a satellite, irrespective of whether there are any active receivers for the given groups. Dynamic forwarding on the satellite interface is required in these cases to control the set of groups that are forwarded from the Gateway to the receivers.

When dynamic forwarding is used, each RCST determines – and, indicates to the NCC – the multicast group traffic that it wishes to be forwarded to it. This information is collected by the Proxy by reception of IGMP/MLD Membership Reports on the LAN interface. The RCST indicates the need to forward traffic over the satellite by sending a group membership message (join) upstream to the Gateway. The protocol used for this control message could be an extension of the protocol used on the LAN interface – i.e. a RCST could proxy IGMP, MLD or PIM, or it could summarize the group membership state into a satellite-specific protocol.

The flow of control information from the set of RCSTs with active receivers allows the Gateway to identify the set of groups that need to be forwarded over the satellite interface, and provides an indication to the NCC for control of the multicast service. This means that satellite capacity is not used for traffic for which there are no active receivers and, that a Gateway router could itself generate an upstream PIM-SM Join message for a requested group to dynamically request the content from a source connected via an upstream-multicast network.

Dynamic forwarding increases the complexity of the multicast service as it necessitates more control functions to realize an effective operational service. For example, the set of groups requested by a RCST needs to be controlled by the NCC to prevent an RCST requesting unauthorized (i.e. outside the SLA for the RCST service) or illegal (i.e. for address ranges that cannot be used) groups. Additional multicast control functions may need to be related to Authentication, Authorization and Accounting, AAA, functions at the Gateway, to implement subscriber control and enable accounting for billing.

The control functions need to be virtualized, if the same multicast traffic is to be forwarded to multiple SVNs. The mappings of content to lower layer addresses may also need to consider address translation when source-specific groups are used with private addresses. An SLA is also needed for each multicast service, to determine the QoS attributes to be used by the Feeder for each group or for all multicast traffic (e.g. peak rate, binding to L2 address and Lower Layer Service on the forward link).

Table 6.1 summarizes distinct forwarding modes discussed so far. Following clauses present more details and examples on these forwarding modes. Table 6.2 summarizes the advantages and disadvantages of dynamic multicast.

**Table 6.1: Multicast modes at satellite and LAN interfaces**

| Mode | Satellite Forwarding | RCST Receiver | LAN Forwarding |
|---|---|---|---|
| Static, passive LAN forwarding | Determined by NCC configuration. Independent of active RCST receivers. | Static by configuration | Independent of the set of active receivers downstream to the RCST. |
| Static, Active LAN forwarding | Determined by NCC configuration. Independent of active receivers. | Forwarding decision made using IGMP/MLD proxy | Only when there are active receivers downstream to the RCST. |
| Dynamic (implies Active forwarding) | Determined by Gateway informed by proxy at RCST, only for 1 or more active RCSTs. | Forwarding decision made using IGMP/MLD proxy | Only when there are active receivers downstream to the RCST. |

**Table 6.2: Advantages and disadvantages of dynamic multicast**

| Advantages | Disadvantages |
|---|---|
| The satellite capacity is only consumed for multicast traffic that is required at the receiver. | Increased inbound control traffic when STs request to receive a multicast group, however the control traffic is usually much less than the total traffic. |
| An RCST can receive arbitrary multicast flows (if permitted by the NCC), including traffic with an IP group destination address not known a priori. | Increased complexity at the Gateway, where multicast control protocols (e.g. PIM) need to be deployed to communicate with upstream multicast networks. |
| It does not require pre-configuration of the Feeder and Gateway router to support specific groups. | The NCC may need to support dynamic construction of the MMT2 and reconfiguration of the Feeder. |
| Provides the service operator with the ability to monitor/charge users for consumed content. | Increased operational complexity in managing and supporting the service. |

## 6.3.1 Static forwarding with Passive mode on the LAN interface

In the passive multicast forwarding mode, the RCST does not participate in IGMP or MLD.

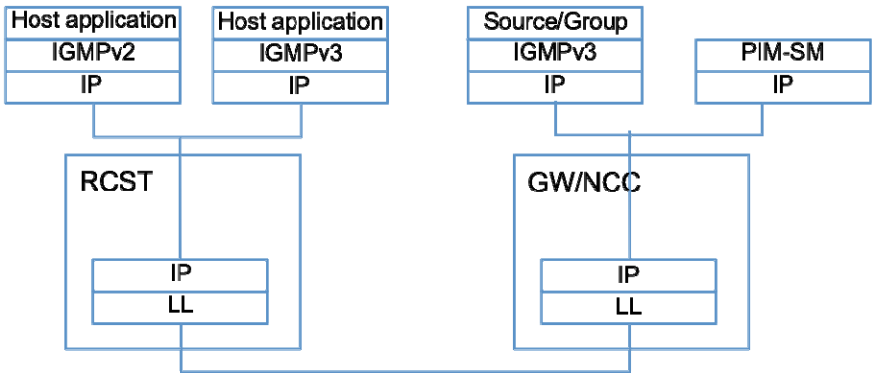


**Figure 6.1: IGMP passive mode for DVB-RCS with static forwarding**

Figure 6.1 shows the passive mode implementation of IGMP for the DVB-RCS2 network. The Gateway has no IGMPv2/v3 stack.

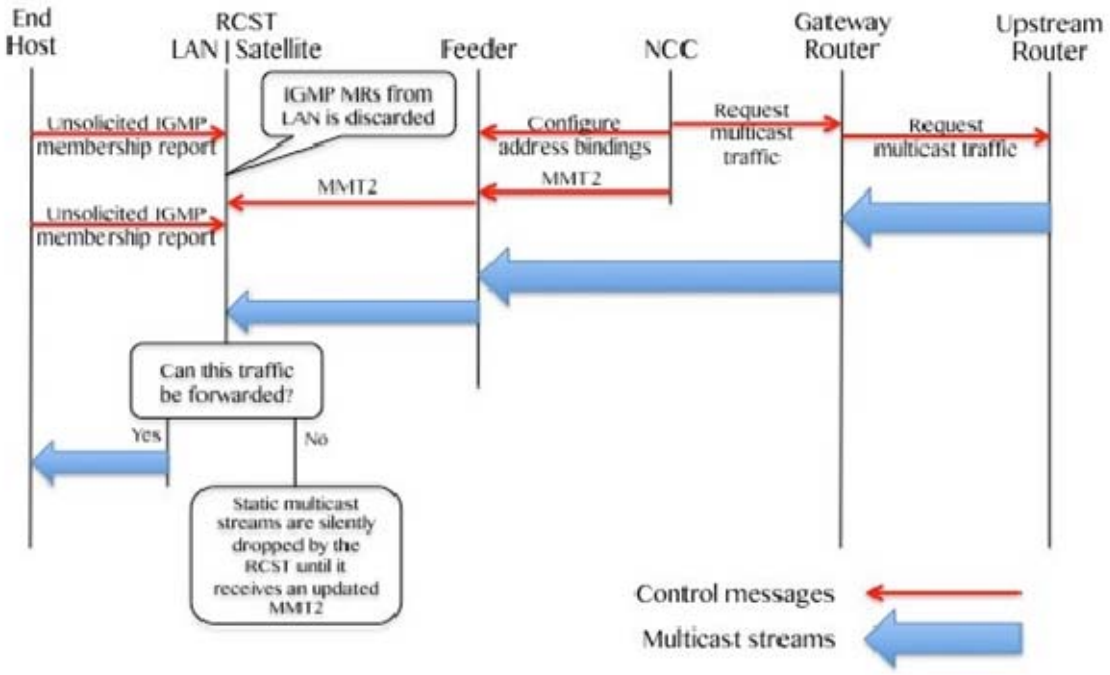


**Figure 6.2: Static multicast with passive forwarding using MMT2**

Figure 6.2 shows a ladder diagram illustrating the delivery of multicast using static forwarding. The NCC configures the Feeder with the mappings for the MAC24 to be used for each active IP multicast address.

In static multicast, the NCC will have previously generated MMT2 entries required and will have configured the Feeder and Gateway Router to forward the multicast stream over the forward link.

The RCST will forward the multicast streams on its LAN interface if it is configured to do so, i.e. if the IP multicast address is in the MFIB. Otherwise, it will silently drop the multicast streams being received. An RCST may inspect the content of the MMT2 to identify all active GSE multicast mappings for the SVN to which it belongs.

## 6.3.2 Static forwarding with Active mode on the LAN interface

Figure 6.3 shows the active mode for IGMP in a DVB-RCS2 network. In active mode, a proxy agent [i.18] is implemented at the RCST. The proxy implements timer values and forwarding rules associated with this active mode.

In active forwarding mode, an IPv4 RCST performs the IGMP router function on its LAN interfaces and the host function of IGMP on its return satellite interface. An RCST should not perform the router portion of IGMP on its return satellite link. However, in the IGMP active mode, the Gateway can be queried on both of its interfaces (forward satellite link as well as the upstream core-network interface).

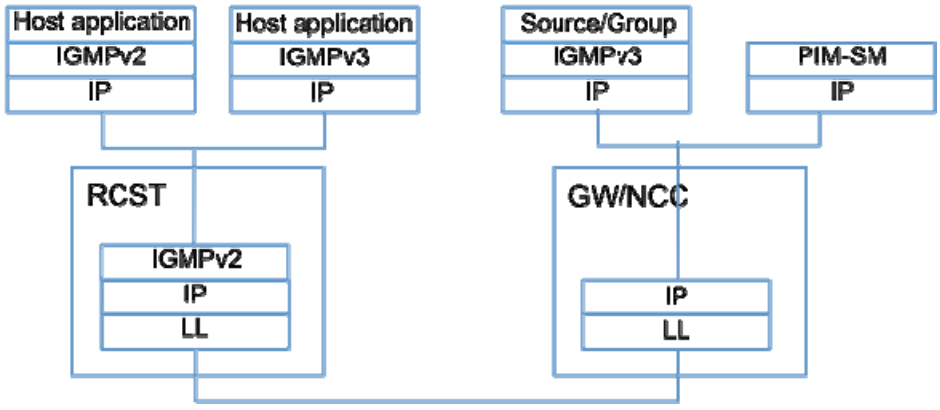


**Figure 6.3: IGMP Active-Mode for DVB-RCS2 with Static Forwarding**

Figure 6.4 provides a ladder diagram showing the delivery of multicast in static configuration mode with active forwarding at the RCST. The only difference in this scenario compared to the previous, is the control of the static multicast streams at the RCST LAN interface.

The RCST Agent collects IGMP/MLD Membership Reports received at the LAN Interface to populate the MFIB with the IP group multicast addresses of the traffic to be received by hosts on the LAN. It uses the information in the MFIB to forward the static multicast streams. This method is the proposed default case for the DVB-RCS2 system.

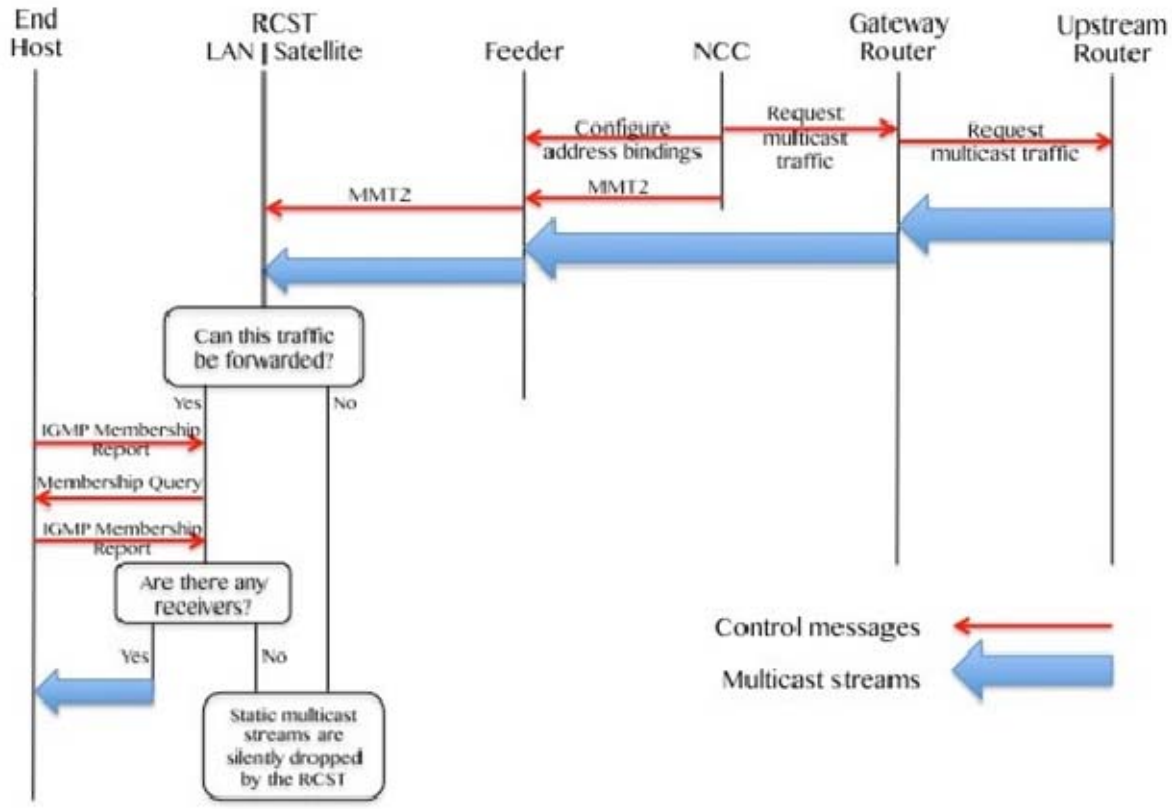


**Figure 6.4: Static multicast with active forwarding using MMT2**

## 6.3.3 Dynamic forwarding with Active mode

When a host sends a multicast membership report to the LAN Interface, the proxy agent at the RCST will forward the request upstream to the Gateway. The Gateway may forward this to the NCC to check the authentication and record the activity.

In dynamic forwarding, the requests from the RCST trigger the Gateway Router to join a specific group on its upstream interface. When necessary, the Feeder is also (re) configured to forward the group, and the NCC may update the MMT2 to reflect any changes made to the Feeder configuration.

Figure 6.5 depicts a ladder diagram for the case of dynamic multicast when the content is already being sent over the satellite air interface (e.g. when a different RCST has requested reception of the same IP multicast group). In this case the NCC has already configured the Feeder and the Gateway Router with entries for the required multicast stream to be forwarded.

Figure 6.6 shows the ladder diagram for dynamic request of new multicast content. In this scenario, the content is not sent via the satellite until requested.

A host sends an IGMP/MLD membership report to the RCST, which then forwards the request over the satellite link to the Gateway.

Since the Gateway does not have an entry for this particular group, it requests authentication from the NCC. When it receives authentication, the Gateway Router send a multicast join request upstream to request the content.

The NCC also performs any required update to the configuration of the Feeder and MMT2. If the update modified the MMT2 address mappings, then an updated MMT2 is required to allow the RCST to receive the multicast content.

Once multicast traffic arrives at the Gateway Router, it forwards the content to the Feeder using the Forward Link. This scenario is different to that of the previous case because it may trigger configuration of the Feeder and the authentication of the request by the NCC, before a user receives the multicast stream.

Pre-configuration of the Feeder can simplify the control interaction. For example, the Feeder could be allowed to forward a block of multicast addresses from the Gateway Router to a specific SVN and could advertise this binding in the MMT2 before it receives any request. This removes the need to reconfigure the feeder each time a request is received for a group address that was mapped.

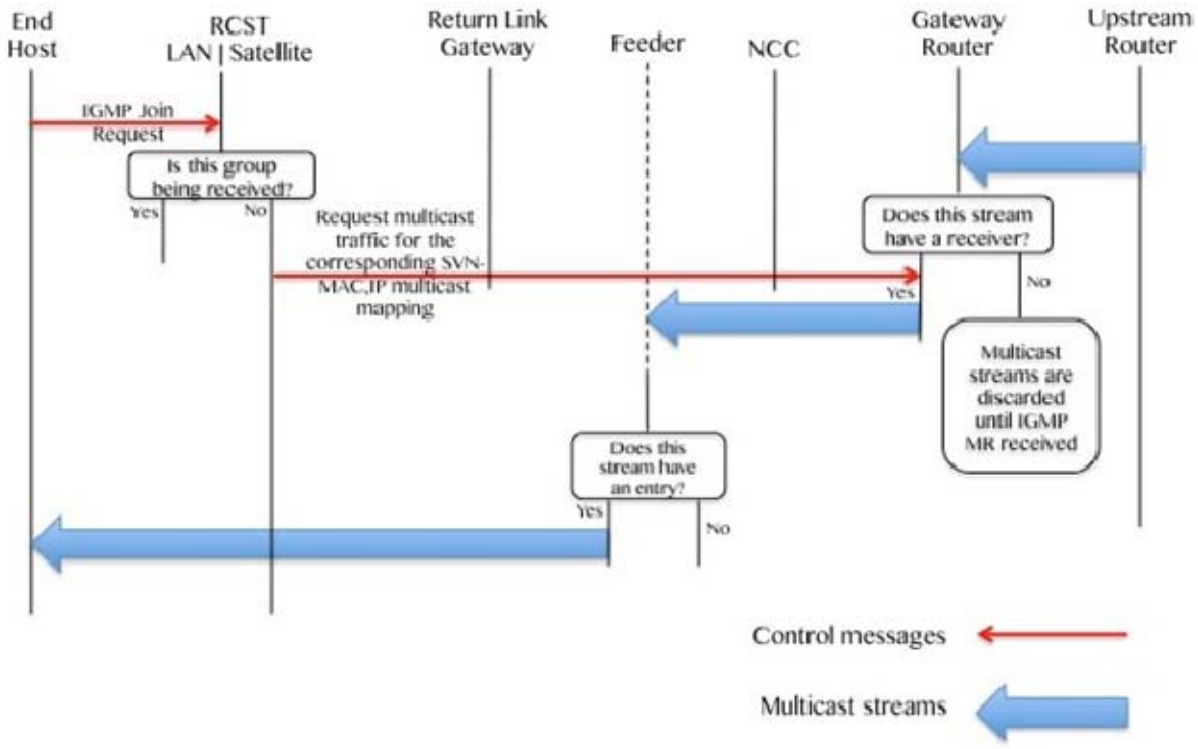


**Figure 6.5: Dynamic multicast request for existing IPv4 multicast content**

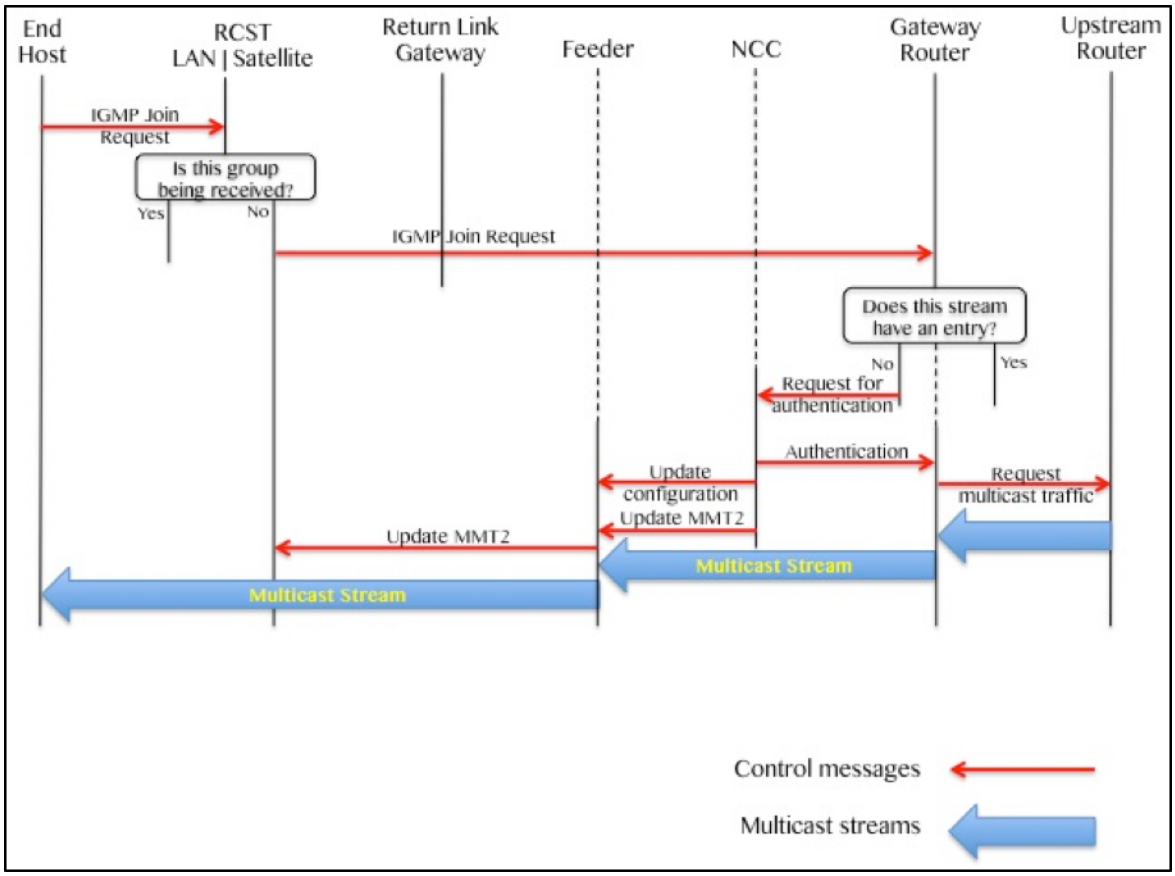


**Figure 6.6: Dynamic request for new IPv4 multicast content**

Figure 6.7 shows the case where the authentication step results in rejection of a request by the NCC.

When the RCST receives a multicast membership report it forwards this to the Gateway, which will then seek authentication from the NCC. If the NCC rejects this request, it will update the Gateway (and the RCST) to ignore all Join requests for the particular multicast group for a specified period. If the RCST receives additional membership reports, requesting to join the group during this interval, they will be silently dropped. It is important to note that the RCST should still forward the membership report for other groups during this interval.

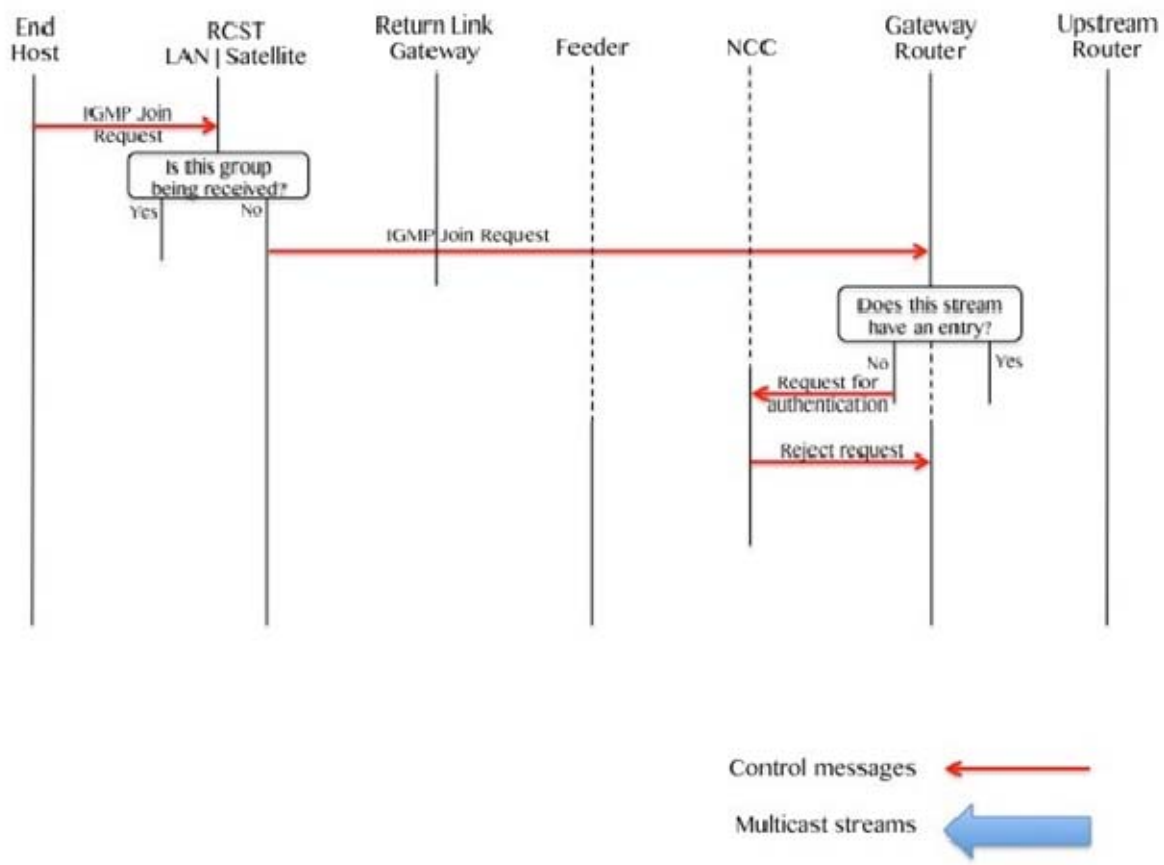


**Figure 6.7: Dynamic request for IPv4 multicast content rejected by the NCC**

### 6.3.4 IP multicast walkthrough in DVB-RCS2

The following entities are required to support a DVB-RCS2 multicast service:

- A multicast-enabled Gateway Router, that may use PIM-SM to request upstream content from the terrestrial network to which it is connected and forward this to the Feeder.
- For the forward link user plane, the NCC should authorize use of forward link satellite capacity for a multicast service by the Feeder and coordinate the use of layer 2 addresses.
- For the forward link control plane, the NCC may also need to generate a set of MMT2 tables for transmission by the Feeder indicating mappings for each SVN that it supports.
- The Gateway receiver, in the case of dynamic multicast should process requests for content (join messages) from an RCST.
- The Feeder should encapsulate and forward multicast flows on the forward link. The feeder also distributes the MMT2 control table to all RCSTs.
- In the user-plane the RCST should enable multicast reception (filtering), forwarding and the processing of multicast address bindings, including parsing of the MMT2.
- In the control plane, the RCST may also need to support an IGMP proxy function and use this to control forwarding and for the dynamic case, return control information to the NCC (the join message).

## 6.4 Encapsulation of IP multicast packets

The functions required for multicast forwarding (see clause 6.3) over a DVB-RCS2 network can result in a range of system designs.

- Static Multicast & Autonomously synthesised MAC24 address: IP traffic at the Feeder is mapped directly to a MAC24 based on the IP group destination address. The NCC enables forwarding of this group by the Feeder. An RCST that is set to receive an IP address (e.g. in the MFIB) maps the address to the corresponding 3B GSE label and then will unconditionally forward all traffic received with the configured IP address to the LAN interface. In static multicast all groups may be forwarded if this is configured.

- Static Multicast & MMT2: The NCC associates a multicast MAC24 for each multicast group. This is used by the Feeder to set the 3B label in an encapsulated multicast packet. The NCC inserts an entry for the IP multicast group address for the corresponding MAC24 in the MMT2. This table is periodically sent by the Feeder to all RCSTs within an SVN. RCSTs are configured to receive IP groups by determining the IP address to be forwarded (e.g. in the MFIB), and then binding this to a MAC24 (e.g. using the MMT2). In the static case, this may default to all advertised multicast content, or could be restricted by local configuration.

- Static Multicast, Active mode LAN interface: As above except the RCST implements a proxy agent. An RCST that receives signalling on its LAN interface adds the corresponding L3 group destination address to its local multicast forwarding state (e.g. in the MFIB). This requires the agent to respond to IGMP/MLD group membership on its LAN interface. Then, it determines the set of 3B MAC24 address required to receive the multicast streams, based on the MMT2. Packets received on one of these MAC24s are then filtered at the IP level based on the local multicast forwarding state, and all traffic that matches is forwarded over the LAN interface.

- Dynamic Multicast: IP traffic at the Feeder is mapped directly to a MAC24 derived from the IP group destination address. The NCC will enable forwarding of this group at the Feeder. All packets of a flow not requested by any RCST will be discarded. The NCC will usually implement a policy to control whether a particular group is enabled and set corresponding QoS parameters for transmission. As in active forwarding, each RCST implements a proxy agent. This controls forwarding to the LAN interface, in addition, the RCST summarizes its local IP forwarding state (from the MFIB) to the NCC to allow it determine which group should be forwarded.

An RCST that implements a Proxy should filter packets initially by MAC24 and finally by IP multicast address, ensuring that only the traffic belonging to the configured group(s) is forwarded (i.e. discard any unwanted traffic that maps to an overlapping MAC24).

The present document does not include consideration of the design of the IGMP/MLD proxy or the PIM-SM router. Neither does it make recommendations on how these protocols should be configured/ adapted to the satellite case.

### 6.4.1 Address mapping for IPv4/IPv6 addresses using the MMT2

Multicast network addresses used in SVNs that belong to different VRF Groups may be identical but correspond to different multicast groups and need to be handled separately (i.e. avoiding address overlap). SVNOs may also need more control of the mapping used in their SVN, e.g. for traffic engineering or to minimize the cost of multicast filtering at RCSTs.

An RCST can identify the MMT2 entries applicable to itself by monitoring the "svn_number" field in the received MMT2 (see [i.3]). The svn_number field is indicated to the RCST during logon.

- An RCST will decode all entries in the MMT2 that match its pre-assigned "svn_number". The MMT2 comprises a set of entries corresponding to blocks of multicast addresses; the number of blocks is specified by the "mapping_sections" field.

- Each block is specified for one GSE Type field. This value is used to differentiate address blocks defined for IPv4, from the ones for IPv6. The address size will be set accordingly.

- For each block, the MMT2 specifies the start and end address. It may also specify values for specific multicast ranges as exceptions to a previously defined multicast address range.

- For each block of address the RCST derives a base address value and length. The IP multicast address is mapped to a 3-Byte MAC24 address, by first identifying a MAC24 base address. The mcast_prefix_length defines the number of significant bits (starting from the most significant bit) that are used from this base address. The remaining bits of the 3-Byte MAC24 address will be mapped directly the low-order bits of the required IP multicast address.

The simplest MMT2 contains one record per SVN that indicates one address block for the MAC24 range (i.e. mask size). In many cases, the mapping from network multicast address to MAC24 can be fairly static, with the mapping revised from time to time, when needing to add/remove a flow, or re-assign an existing flow to a new multicast address. In DVB-RCS2, the table is created at the NCC and used by the Feeder. Procedures may be defined to automate creation of the MMT2 entries, for instance, to create a mapping based on reception of a dynamic multicast join request from a SVNO (e.g. generated as a result of arriving multicast traffic or reception of IGMP/MLD messages).

The MMT2 structure allows an RCST to map an IP group destination address to any SVN. In normal operation, an SVNO is expected to use a part of the address allocation for the SVN to support the multicast services. In this case the base address retrieved from the MMT2 for a multicast group will likely be a subset of the SVN address range used by the RCST for its unicast service.

The MAC24_base address and mcast_prefix_length can be configured to suit the various multicast scope required by the satellite virtual network operators or the satellite network operator. Note that this scheme is not designed to allow a change of the mcast_prefix_length while RCSTs are logged onto the system. The recommended procedure is therefore to reset the mask length and, for users to then force RCSTs to logon to the updated system using the new mask length. This design decision is justified in that it is not expected that reconfiguration will be required.

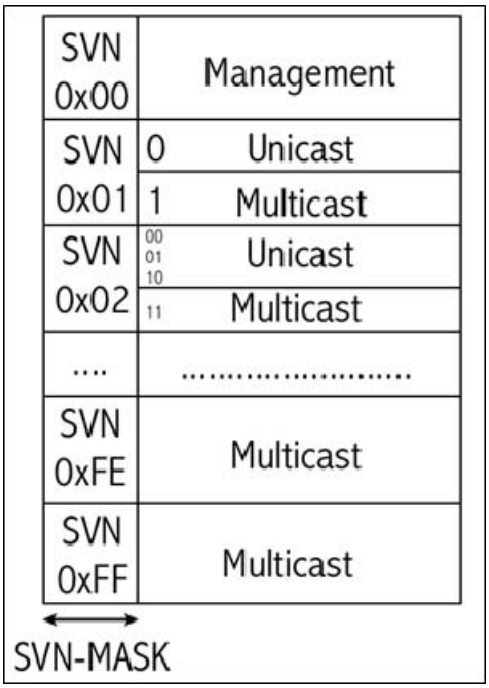


**Figure 6.8: Example MAC24 address allocation to SVNs**

The MMT2 uses a 16-bit encoding to represent the SVN number used by a Satellite Virtual Network Operator. This identifies which systems read the MMT2. Note that the SVN-ID is encoded as a 16 bit number, corresponding to the two most significant bytes of the lowest used MAC24 in the allocated block. This is the form used in the MMT2 to advertise the multicast prefixes, and is used so that the network can accommodate multiple sizes of svn_mask. This allows an SVNO to subdivide a single Operator Virtual Network (OVN) allocation from an SNO to realize multiple SVNs within its assigned OVN. Hence, if an SVNO is allocated 0x0100000-0x01FFFF, the MMT2 entry may be 0x0100, but if the operator chooses to subdivide this OVN allocation into two, then he could separately generate a MMT2 for 0x01800 and 0x01000, since RCSTs always the know the SVN_mask length, they know whether the OVN has been subdivided in this way.

SNOs will allocate SVNOs with "*svn_numbers"*, which are then used by them to assign address ranges amongst their users (RCSTs). During startup RCSTs are informed of their "*svn_number"*, SVN-MASK length and assigned a unicast MAC24 by the SVNO.

Each SVNO can allocate parts of the allocated address as either unicast, multicast or reserved for future use. The allocation to be used by a particular SVN is notified in the Multicast Mapping Table 2 (MMT2). The MMT2 contains a list of mappings for multicast addresses for each "*svn_number"*, i.e. for each SVN.

An RCST receives the MMT2 using multicast. It examines "svn_number" of received messages and only accepts messages where the number matches a svn_number value assigned to one of the RCST interfaces.

## 6.4.2 Mapping for IPv4/IPv6 addresses to the same MAC24 prefix

The MMT2 mapping should by default use different MAC24 prefixes for IPv4 and IPv6 traffic. This use resembles the use in Ethernet in GSE, where overlapping between address ranges do not exist, because IPv4 and IPv6 traffic is assigned to a different Organizationally Unique ID, OUI.

In some cases, SVN address space can be conserved by mapping the two sets of IP addresses to the same MAC24 base address. However, this can also result in overlap between IPv4 and IPv6 multicast. This overlap could have unwanted side-effects.

One scenario where this separation is desirable would be when a content provider maps content to both IPv4 and IPv6 and both use a format where the least significant part is the same. When an RCST dynamically registers for IPv4 membership join, it opens the L2 filter to receive the content. A subsequent request by an RCST that desires IPv6 content would then lead to this additional traffic being sent and would be passed by the existing L2 filter, requiring protocol filtering at the higher layer. While this filtering will only forward the requested IPv4 traffic, the IPv6 group still contributes to additional unnecessary processing cost at both L2 and L3.

## 6.4.3 Aliasing for IPv4/IPv6 addresses using the MMT2

An SVNO may decide through bilateral agreement with another SVNO or the SNO to use a MAC24 multicast prefix that lies outside the unicast address range that it uses. In this case, the base address retrieved from the MMT2 for a multicast group will belong to a different SVN address range. This effectively allows one operator to signal use of another block of addresses. This method could be used to group together multicast services for several SVNs and may eliminate the otherwise need to replicate common multicast streams for different services. Care needs to be exercised in using this method so that the addresses remain in scope (i.e. the aliased address has the same meaning in the SVN for which it is to be received). Hence, the MMT2 may be used to support a network group that is accessible from more than one SVN and is mapped to a common MAC24. The SNO/SVNO is responsible for such system-wide co-ordination of the use of MAC24 labels.

## 6.4.4 Example address mappings using MMT2

### 6.4.4.1 Simple MAC24 mapping for multicast address blocks

**Table 6.3: SVN 0x01 Example address block allocation**

| SVN 0x01 Address Block (mcast_prefix_length = 9 bits) | Allocation |
|---|---|
| 0x010000 – 0x017FFF | Unicast |
| 0x018000 – 0x01FFFF | Multicast |

The MMT2 defines a multicast SVN block using the "mcast_prefix_length", adding a single bit to the existing 8-bit (in this example) SVN-MASK (i.e. 9 bits). In Table 6.3, a SVNO that has been assigned 0x0100 and has divided its address block of 64K (approx.) address into two address blocks of 32K addresses for unicast and 32K for multicast. The SVNO therefore assigns a "mcast_prefix_length" of 9 bits. Table 6.3 shows the allocation of SVN 0x0100 with a different division of the address block. The Feeder will be configured with this allocation.

**Table 6.4: MMT2 example for SVN 0x0100**

| MMT2 | | |
|---|---|---|
| **Field** | **Value** | **Length** |
| svn_number | 0x0100 | 16 bits |
| pt_count | 1 | 8 bits |
| protocol_type | 0x800 | 16 bits |
| address_size | 0x04 (4 Bytes) | 8 bits |
| mapping _section | 1 | 8 bits |
| inclusive_start | 224.0.0.1 | 32 bits (4B) |
| inclusive_end | 239.255.255.255 | 32 bits (4B) |
| exclusions | 0 | 8 bits |
| mac24_base | 018000 | 24 bits |
| mcast_prefix_length | 01001 (9 bits) | 5 bits |

The MMT2 shown in Table 6.4 is an example of the most basic scenario of multicast address mapping to a MAC24. In this example, the SVNO has mapped the complete multicast address range to a block of addresses within its address range. Table 6.4 illustrates the MMT2 used for SVN 0x0100. In this scenario, the Feeder generates the MMT2 during initial configuration and no further updates will be necessary to support requests for dynamic multicast streams because the entire multicast address block has been assigned.

### 6.4.4.2 Dynamic MAC24 mapping for multicast address blocks

This example provides a use case for a dynamic multicast stream, where the Feeder will map the IP multicast address to an assigned MAC24 label in the range 0x02C000 – 0x02FFFF (Table 6.5). The value of mcast_prefix_length changes to 10 bits in this use case. The Feeder will transmit the MMT2 for the SVN (0x0200) to inform the RCST of mappings for the multicast stream. Table 6.6 shows the MMT2 for SVN 0x0200 in which the RCST is informed of the IP multicast mapping to a MAC24.

**Table 6.5: SVN 0x0200 Example address block allocation**

| SVN 0x02 Address Block (mcast_prefix_length = 10 bits) | Allocation |
|---|---|
| 0x020000 – 0x023FFF | Unicast |
| 0x024000 – 0x027FFF | Reserved for future use with Unicast or Dynamic Multicast |
| 0x028000 – 0x02BFFF | |
| 0x02C000 – 0x02FFFF | Multicast |

**Table 6.6: MMT2 example for SVN 0x0200**

| MMT2 | | |
|---|---|---|
| **Field** | **Value** | **Length** |
| svn_number | 0x0200 | 16 bits |
| pt_count | 1 | 8 bits |
| protocol_type | 0x800 | 16 bits |
| address_size | 0x04 (4 Bytes) | 8 bits |
| mapping _section | 1 | 8 bits |
| inclusive_start | 224.0.0.1 | 32 bits (4B) |
| inclusive_end | 239.255.255.255 | 32 bits (4B) |
| exclusions | 0 | 8 bits |
| mac24_base | 02C000 | 24 bits |
| mcast_prefix_length | 01010 (10 bits) | 5 bits |

#### 6.4.4.3 MAC24 mapping using the "exclusions" field

The MMT2 syntax permits flexibility for mapping an IP multicast address to MAC24 labels. The "exclusion" field allows a SVNO to specify a different MAC24 mapping behavior for different IP multicast address ranges and assign these to MAC24 labels, i.e. dynamic allocation. Table 6.7 shows an MMT2 for SVN 0x0300, which has the same format of address block allocation as SVN 0x0200 (shown in Table 6.5). In this SVN (0x0300), the SVNO has used the "exclusion" field to exclude the SSM address range (as shown in Table 6.7). A use case may be to support dynamic multicast, where the MAC24 is assigned when multicast forwarding is setup. In this example, the Feeder dynamically assigns one SSM address with an MAC24 address from the unallocated range (shown in Table 6.7) with the use of a second "inclusion_start" and "inclusion_stop" section within the MMT2. A second "rcs_mac_base" field along with the "mcast_prefix_length" is used to assign an MAC24 labels from an address block different to that allocated in the first section.

**Table 6.7: MMT2 example for SVN 0x0300**

| MMT2 | | |
|---|---|---|
| **Field** | **Value** | **Length** |
| svn_number | 0x0300 | 16 bits |
| pt_count | 1 | 8 bits |
| protocol_type | 0x800 | 16 bits |
| address_size | 0x04 (4 Bytes) | 8 bits |
| mapping _section | 2 | 8 bits |
| inclusive_start | 224.0.0.1 | 32 bits (4B) |
| inclusive_end | 239.255.255.0 | 32 bits (4B) |
| exclusions | 1 | 8 bits |
| exclusion_start | 232.0.0.0 | 32 bits (4B) |
| exclusion_stop | 232.255.255.255 | 32 bits (4B) |
| mac24_base | 03C0000 | 24 bits |
| mcast_prefix_length | 01010 (10 bits) | 5 bits |
| inclusive_start | 232.0.0.1 | 32 bits (4B) |
| inclusive_end | 232.0.0.1 | 32 bits (4B) |
| exclusions | 0 | 8 bits |
| mac24_base | 038000 | 24 bits |
| mcast_prefix_length | 11000 (24 bits) | 5 bits |

In the example shown in Table 6.7, allocation of an additional SSM address to a MAC24 label may be performed by the Feeder dynamically upon a new request and then signalling by the additional inclusion section of the MMT2 sent to the RCSTs.

### 6.4.5 Address mapping for non-IPv4 addresses

Many protocols also use L2 multicast apart from IPv4 and IPv6. A method may be provided to support non-IP multicast. This could be done by mapping a non-IP multicast L3 address to a L2 address, or by mapping between the LAN MAC address and the MAC24 label. The MMT2 supports non-IP multicast services, with the mappings identified by the use of their allocated protocol type field within the table.

Dynamic methods are not specified and will rely on definition of an agent and a satellite control protocol between the agent and the NCC (e.g. an adapted IGMP/MLD or PIM-SM stack).

### 6.4.6 Address-specific issues

The IETF specifies the use of IP multicast addresses. The currently allocated set of multicast addresses by the Internet Assigned Numbers Agency (IANA) was summarized in [i.20]. This also provides general guidance on the use of the multicast address space and defines the procedures for address allocation within the multicast address blocks. With the exception of some reserved addresses the allocation of an IPv4 multicast addresses to groups is dynamic. Well-known multicast sources may be allocated a fixed and advertised multicast address.

Specific multicast addresses have been statically allocated to certain roles, especially when these relate to specific protocols.

Figure 6.10 shows the IANA-allocated multicast address blocks from the perspective of a satellite network. IP multicast address allocation within a satellite network has to be carefully assigned by the SNOs and SVNOs. The assignment of multicast domains and RPs has to be performed by the SNO as well.

In Figure 6.9, SVN 0xFF is dedicated to globally-assigned multicast using the GLOP Block and the ADHOC Block III. These addresses are globally unique multicast addresses assigned by network operators. This will be used by the SNO to avoid simulcast of the globally unique multicast amongst the SVNOs in the SNO network.

In this example, the SVN 0x0100 is also allocated a shared SVN, with MAC24 block dedicated for all other multicast traffic. This is assigned the start address of 0xFE. This can be used by the SVNO for shared multicast distribution. This avoids the need to simulcast the same content in different SVNs.

The local network control block of multicast addresses (224.0.0.0 – 224.0.0.255) needs to be assigned to a preconfigured block of multicast MAC24 labels for each SVNO. These traffic flows should be independent per each IP network.

Allocations have to performed to ensure that multiple SVNs are mapped to the appropriate domain, i.e. local network control packets has to be delivered between multiple SVNs if the SVNs are in the same domain.

Actual allocations do not need to be for entire multicast ranges, and do not need to use a SVN-MASK of 8 bits, as used in these examples shown in Tables 6.11, 6.12 and 6.13.

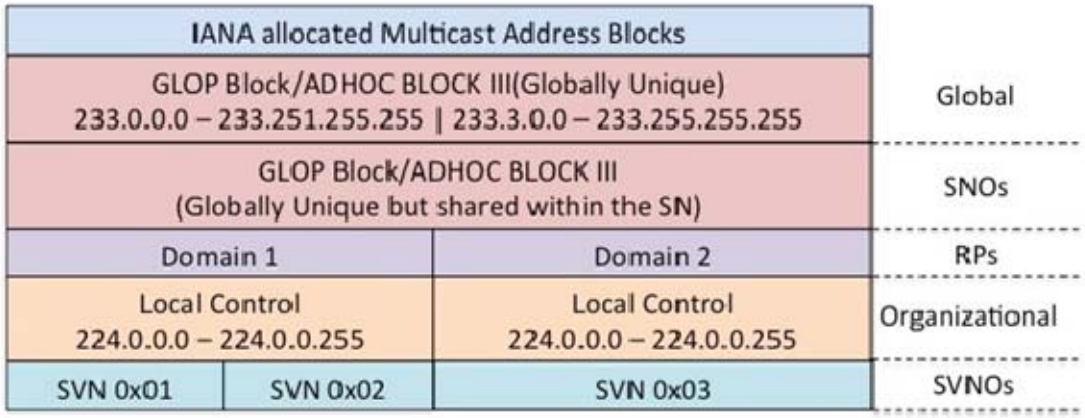


**Figure 6.9: IANA allocated Multicast Address Blocks and their example mapping to SVNO/SVNs**

**Table 6.8: MMT2 example for SVN 0x0100 using shared transmission**

| MMT2 | | |
|---|---|---|
| **Field** | **Value** | **Length** |
| svn_number | 0x0100 | 16 bits |
| pt_count | 1 | 8 bits |
| protocol_type | 0x800 | 16 bits |
| address_size | 0x04 (4 Bytes) | 8 bits |
| mapping _section | 3 | 8 bits |
| inclusive_start | 224.0.0.1 | 32 bits (4B) |
| inclusive_end | 224.0.0.255 | 32 bits (4B) |
| Exclusions | 0 | 8 bits |
| mac24_base | 017F00 | 24 bits |
| mcast_prefix_length | 10000 (16 bits) | 5 bits |
| inclusive_start | 224.0.1.0 | 32 bits (4B) |
| inclusive_end | 232.255.255.255 | 32 bits (4B) |
| Exclusions | 0 | 8 bits |
| mac24_base | FE0000 | 24 bits |
| mcast_prefix_length | 01000 (8 bits) | 5 bits |
| inclusive_start | 233.0.0.1 | 32 bits (4B) |
| inclusive_end | 233.255.255.255 | 32 bits (4B) |
| Exclusions | 0 | 8 bits |
| mac24_base | FF0000 | 24 bits |
| mcast_prefix_length | 01000 (8 bits) | 5 bits |

**Table 6.9: MMT2 example for SVN 0x0200 using shared global allocations and SVN-local address for other multicast traffic**

| MMT2 | | |
|---|---|---|
| **Field** | **Value** | **Length** |
| svn_number | 0x0200 | 16 bits |
| pt_count | 1 | 8 bits |
| protocol_type | 0x800 | 16 bits |
| address_size | 0x04 (4 Bytes) | 8 bits |
| mapping _section | 3 | 8 bits |
| inclusive_start | 224.0.0.1 | 32 bits (4B) |
| inclusive_end | 224.0.0.255 | 32 bits (4B) |
| exclusions | 0 | 8 bits |
| rcs_mac_base | 027F00 | 24 bits |
| mcast_prefix_length | 10000 (16 bits) | 5 bits |
| inclusive_start | 224.0.1.0 | 32 bits (4B) |
| inclusive_end | 232.255.255.255 | 32 bits (4B) |
| exclusions | 0 | 8 bits |
| rcs_mac_base | FE0000 | 24 bits |
| mcast_prefix_length | 01010 (10 bits) | 5 bits |
| inclusive_start | 233.0.0.1 | 32 bits (4B) |
| inclusive_end | 233.255.255.255 | 32 bits (4B) |
| exclusions | 0 | 8 bits |
| rcs_mac_base | FF0000 | 24 bits |
| mcast_prefix_length | 01000 (8 bits) | 5 bits |

**Table 6.10: MMT2 example for SVN 0x0300 using shared global allocations**

| MMT2 | | |
|---|---|---|
| **Field** | **Value** | **Length** |
| svn_number | 0x0300 | 16 bits |
| pt_count | 1 | 8 bits |
| protocol_type | 0x800 | 16 bits |
| address_size | 0x04 (4Bytes) | 8 bits |
| mapping _section | 3 | 8 bits |
| inclusive_start | 224.0.0.1 | 32 bits (4B) |
| inclusive_end | 224.0.0.255 | 32 bits (4B) |
| exclusions | 0 | 8 bits |
| mac24_base | 037F00 | 24 bits |
| mcast_prefix_length | 10000 (16 bits) | 5 bits |
| inclusive_start | 224.0.1.0 | 32 bits (4B) |
| inclusive_end | 232.255.255.255 | 32 bits (4B) |
| exclusions | 0 | 8 bits |
| mac24_base | FE0000 | 24 bits |
| mcast_prefix_length | 01000 (8 bits) | 5 bits |
| inclusive_start | 233.0.0.1 | 32 bits (4B) |
| inclusive_end | 233.255.255.255 | 32 bits (4B) |
| exclusions | 0 | 8 bits |
| mac24_base | FF0000 | 24 bits |
| mcast_prefix_length | 01000 (8 bits) | 5 bits |

### 6.4.7 Source-specific multicast support with MMT2

Currently, there is no defined syntax to support the MMT2 describing the source address for a group.

## 6.5 Multicast management for DVB-RCS2

The control for delivery of IPv4 multicast in static mode with active forwarding at the RCST is provided using an IGMP MIB at the RCST. This is the default case for DVB-RCS2. A candidate IGMP MIB for a RCS2 network is defined in Tables 6.11 and 6.12.

The RCST should show, for debugging purposes, information about the multicast sessions it is subscribed, that is, host information to be reported to the querier when using IGMP. Management of multicast requires visibility of the active address mappings and also the filters used at an RCST, since the Service Provider will need to confirm the set of active multicast groups, the mappings to L2 and the status of group membership subscription via PIM, IGMP and/or MLD. This allows an operator to determine whether a multicast outage is due to a L3 routing/RPF issue, a L2 problem, or an upstream network problem. Traffic statistics can show information about the forwarded packets over the satellite and LAN interfaces.

A multicast forwarding table defines the configuration items in the RCST, in case they are not indicated upon logon procedure in the HLS support descriptor. This applies to static modes, where the installer or the SVNO configures the forwarding mode of the RCST.

### 6.5.1 Multicast configuration and monitoring in RCST MIB

The MIB for DVB-RCS may be organized in two tables: the Interface and Cache tables (Tables 6.11 and 6.12). The IGMP Interface table contains entries for each interface that supports IGMP on a device. For the Gateway, this includes the Core-Network and satellite interfaces, while for the RCST, the satellite and LAN interfaces. The IGMP Cache table contains one row per each IP Multicast Group for which there are active members on a given interface. Active membership should only exist on the RCST LAN interface. However, active membership may exist on both the network side and satellite interfaces of the Gateway.

**Table 6.11: IGMP interface table in RCST**

| | RCST | |
|---|---|---|
| | Active | |
| **igmpInterfaceTable** | **Upstream-Network Side** | **RCST LAN network** |
| igmpInterfaceIfIndex | Not-accessible, assigned interface number | Not-accessible, assigned interface number. |
| igmpInterfaceQueryInterval | Read-only, the RCST should not transmit queries upstream, return 0. | Read-create, min = 0, max = (2^32-1), default = 125 |
| igmpInterfaceStatus | Should be enabled on both interface for all DVB-RCS RCST interfaces | |
| igmpInterfaceVersion | Should be version 2 for all DVB-RCS RCST interfaces | |
| igmpInterfaceQuerier | Read-only, Should be the address of an upstream IGMP Querier device for both active and passive RCSTs. | Read-only, active RCSTs may report it as the satellite interface value. However, active RCSTs that participate in IGMP Querier negotiation on the RCST LAN interface may report it as a different RCST LAN device. |
| igmpInterfaceQueryMaxResponseTime | n/a, read-only, return value of 0 | Read-only, value derived from observation of queries received from an upstream querier. |
| igmpInterfaceQuerierUpTime | Read-only | |
| igmpInterfaceQuerierExpiryTime | n/a, read-only, return 0 | Read-only, RCST may only be the querier on the RCST LAN |
| igmpInterfaceVersion1QuerierTimer | Read-only | |
| igmpInerfaceWrongVersionQuerier | Read-only, the number of non-v2 queries received on this interface. | |
| igmpInterfaceJoins | n/a, read only, return 0 | Read-only, group membership defined to only exist on the RCST LAN. |
| igmpInterfaceProxyIfIndex | Read-only, return 0 | Read-only, return a ifIndex for satellite-link interface |
| igmpInterfaceGroups | n/a, read only, return 0 | Read-only, group membership is defined to exist on the RCST LAN interface. |
| igmpInterfaceRobustness | Read-create, min = 1, max = (2^32 – 1), default = 2. | |
| igmpInterfaceLastMembQueryIntv1 | n/a, read-only, return 0 | Read-create, min = 0, max = 255, default = 100 |

**Table 6.12: IGMP cache table at RCST**

| **igmpCacheTable** | | | |
|---|---|---|---|
| igmpCacheAddress | Not-accessible (index), report the address of active IP Multicast on the RCST LAN interface. | | |
| igmpCacheIfIndex | Should only apt to RCST LAN interface. | | |
| igmpCacheSelf | Read-create, implementation specific. If RCST configured to be member of group, then membership reports are sent with the RCST's IP address but SHOULD ONLY be sent in proxy for active sessions. If the RCST is not configured to be a member, then the source IP address of membership reports SHOULD be set to the current value of the igmpCacheLastReporter address. | | |
| igmpCacheLastReporter | Should only apply to last reporter on RCST LAN interface. | | |
| igmpCacheUpTime | Read-only, Should only apply to duration of membership on RCST LAN interface. | | |
| igmpCacheExpiryTime | Read-only, Should only apply to duration of membership on RCST LAN interface. | | |
| igmpCacheStatus | Read-create. Should only apply to membership on RCST LAN interface. | | |
| igmpCacheVersion1HostTimer | n/a, read-only, return 0 | Read-only | |

### 6.5.2 Multicast forwarding management

New *multicastFiltersTable* (see Table 6.13) is needed for including the Layer 2 filters that will be used by the ST to receive the multicast streams and forward them to the user LANs/VLANs. This table is unique for all the SVNs present in the ST. According to the value of parameter *vrfMulticastMappingMethod*, the mechanism for the population of this table differs. In case the MMT2 method is used, the table is automatically composed and updated by the ST upon MMT2 decoding. Each IP multicast flow should be delivered to the corresponding SVN interface, only when their respective membership group is active. The SVN interface can be deducted from the respective SVN number.

**Table 6.13: Multicast Filter MIB table**

| Element | Range | Description in HLS | Changes for HLS |
|---|---|---|---|
| multicastFiltersTable | SEQUENCE OF multicastFiltersTableEntry | - | New table |
| multicastFiltersTableEntry | SEQUENCE { multicastFilterIndex, multicastFilterSVNnumber, multicastFilterRCSMAC, multicastFilterInclStart, multicastFilterInclEnd, multicastFilterExclStart, multicastFilterExclEnd, multicastFilterStatusRow } | - | |
| multicastFilterIndex | INTEGER | - | Table index of the multicast entry. |
| multicastFilterSVNnumber | INTEGER | | SVN where the multicast flow should be delivered. This is a link to the virtual interface in *vrfGroupTable*. |
| multicastFilterRCSMAC | OCTET STRING | - | MAC24 of the multicast group. |
| multicastFilterInclStart | InetAddress | - | First multicast IP address included in the range. |
| multicastFilterInclEnd | InetAddress | - | Last multicast IP address included in the range. |
| multicastFilterExclStart | InetAddress | - | First multicast address excluded in the range. |
| multicastFilterExclEnd | InetAddress | - | Last multicast address excluded in the range. |
| multicastFilterStatusRow | Row Status | - | The row status, used according to row creation and removal conventions. A row entry cannot be modified when the status is marked as active(1). A row can be created either by createAndGo and automatically change to active state or createAndWait to add more parameters before becoming active. |

### 6.5.3 Multicast statistics

The statistics for transmitted and received multicast packets can be obtained from the interfaces MIB group, in parameters ifInNUcastPkts and ifOutNUcastPkts, by locating the corresponding SVN interface (identified by parameters ifIndex and ifPhysAddress).

# 7 QoS support

This clause presents an overview of the Differentiated Services (DiffServ) Quality of Service (QoS) model for DVB-RCS2 transmission systems together with implementation guidelines regarding its use. It briefly describes the current IP DiffServ model in terrestrial networks, and describes this model for DVB-RCS2 systems. Finally, it provides examples of QoS configuration for a range of terminal profiles.

The Higher Layers (HL) support at the RCST contains the relevant components to implement QoS support on the Return Link (RL). This includes traffic classification, policing functions, and scheduling according to the HL service associated with traffic flows.

The RCST QoS model in DVB-RCS2 HLS [i.4] aims to satisfy the capacity requirements for different users and services. The four basic components of the QoS model in DVB-RCS2 HLS are:

- RCS2 satellite terminal, RCST.
- Network Control Centre (NCC) and Network Management Centre (NMC). The NCC controls the interactive network (control plane). The NCC is in charge of element and network management functions (in the management plane).
- RCS2 Gateway (GW).
- Operations Support System (OSS).

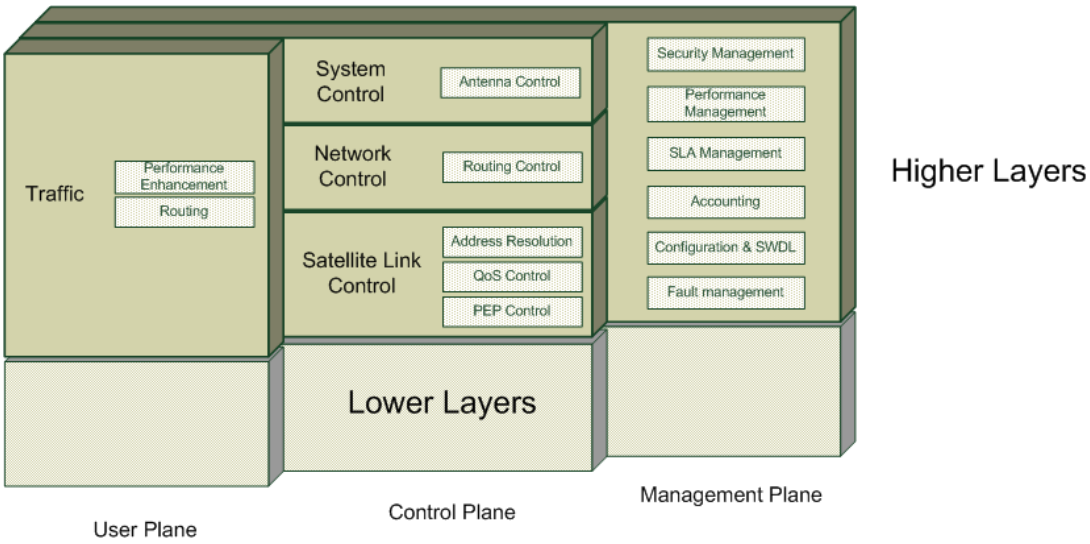


**Figure 7.1: Planes in the higher layers of DVB-RCS2 HLS standard**

The RCST SW may be customized for a given terminal profile. Although, it may also be compatible across multiple terminal profiles. The main DVB-RCS2 terminal profiles are:

- Consumer / SOHO
- Institutional / Corporate
- Backhauling
- Multi-dwelling
- SCADA

Table 7.1 provides an example of the functional QoS requirements that can be provided for different terminal profiles:

**Table 7.1: QoS requirements per terminal profile**

| | Consumer/ SOHO | Institutional/ Corporate | Backhauling | Multi-dwelling | SCADA |
|---|---|---|---|---|---|
| Number of IP QoS classes for HL Service | Up to 5 | 5-7 | 1-3 (GSM)<br>3-5 (maritime, in-flight, on train) | 5 | 1-2 |
| Traffic profiles for transmission and reception | Asymmetric (consumer)<br>A/Symmetric (SOHO) | Symmetric<br>Low-latency | Asymmetric | Fair sharing between users | Asymmetric<br>Bursty traffic |
| User Services | Internet, VoIP, VPN, P2P, gaming, streaming | Corporate<br>Military<br>Surveillance<br>Disaster Relief | Internet in-flight.<br>GSM<br>Satellite/LTE | Internet access, VoIP | Monitoring of real time industrial processes.<br>Environmental monitor |

The DiffServ model [i.21], [i.22] and [i.23] defines an IP QoS architecture based on packet marking. In this model, policy-based management mechanisms are used for prioritizing network resources to meet the requirements of specific traffic types on a per hop basis. No explicit signalling is used to communicate with DiffServ routers, instead a set of Traffic Classifiers (TCs) are used to assign flows to one of a set of pre-defined Behaviour Aggregates (BAs). The classification is performed by inspecting packet header fields, such as IP addresses, ports, and the Differentiated Service Code Point (DSCP) [i.21].

The DiffServ model defines consistent QoS operation within the routers that form a part of the network called the DiffServ domain. The domain consists of a contiguous set of routers operating with a common set of service provisioning policies.

It is common practice to provide traffic conditioning, including admission control, shaping and policing at the edge of a DiffServ domain [i.24]. The DiffServ framework for policy-based admission control [i.25] describes the various components that participate in policy decision-making (i.e. Policy Decision Point, PDP, or Policy Enforcement Point, PEP). Traffic conditioning of admitted traffic may be performed using "meters" to measure the properties of each BA [i.23], [i.26] and [i.27] against a traffic class (or traffic specification). A Policy Enforcement Point may police the PDUs from non-conformant flows (i.e. These may be marked, dropped or shaped). The treatment of the traffic forming a BA is characterized by a Per Hop Behaviour (PHB) [i.21] and [i.24].

Within a DiffServ domain, the network operators may choose to support any combination of standard or operator-specific PHBs. The current set of standard PHBs defined by the IETF is:

- Expedited Forwarding (EF) [i.28] and [i.29]
- Assured Forwarding (AF) [i.30]
- Default (Best Effort) [i.21]

Each standard PHB has been assigned a standard DSCP. EF has been assigned DSCP 46 and BE has been assigned 0. The DSCPs assigned for the AF PHB group are given in Figure 7.2.

Assured Forwarding (AF) Behavior Group

| | Class 1 | Class 2 | Class 3 | Class 4 |
|---|---|---|---|---|
| Low Drop | AF11 (DSCP 10) | AF21 (DSCP 18) | AF31 (DSCP 26) | AF41 (DSCP 34) |
| Med Drop | AF12 (DSCP 12) | AF22 (DSCP 20) | AF32 (DSCP 28) | AF42 (DSCP 36) |
| High Drop | AF13 (DSCP 14) | AF23 (DSCP 22) | AF33 (DSCP 30) | AF43 (DSCP 38) |

**Figure 7.2: PHB assignment to DSCP for AF PHB group**

Other PHBs may also be standardized and vendors/operators may also introduce their own PHBs.

# 7.1 QoS Model in DVB-RCS2

The DVB-RCS2 specification adopts the DiffServ model. Each RCST uses a set of TCs specified in a configured policy(class map) to map packets received on the LAN Interface to a specific HL Service (see clause 7.1.3.1). The set of classified packets handled by a HL Service form a BA. Traffic conditioning of admitted traffic may optionally be performed.

All packets assigned to a BA receive the same treatment (i.e. the same HL Service, that is, they are assigned to the same queuing and IP scheduler treatment). This treatment is characterized by a PHB. Within the RCST, each HL Service is mapped to a Lower Layer Service (LL Service) and a set of Request Classes (RCs) in the control plane. The following table provides examples of the relationships between a set of IP service classes and applications, based on IETF recommendations [i.31].

**Table 7.2: PHB – Example application mapping**

| Service Class Name | DSCP Name | Application Examples |
|---|---|---|
| Network Control | CS6 | Network routing |
| Telephony | EF | IP Telephony bearer |
| Signalling | CS5 | Telephony signalling |
| Multimedia Conferencing | AF41, AF42, AF43 | H.323/V2 Video conferencing |
| Real Time Interactive | CS 4 | Video conferencing and Interactive gaming |
| Multimedia Streaming | AF31, AF32, AF33 | Streaming video and audio on demand |
| Broadcast Video | CS3 | Broadcast TV & live events |
| Low-Latency Data | AF21, AF22, AF23 | Client/Server transactions Web-based ordering |
| OAM | CS2 | OAM&P |
| High-Throughput Data | AF11, AF12, AF13 | Store and forward applications |
| Standard | DF (CS0) | Undifferentiated applications |
| Low-Priority | CS1 | Any flow that has no BW assurance |

Considering Table 7.2 and the terminal profile requirements (in Table 7.1), Table 7.3 can be used to identify a set of BAs that would be appropriate for a specific deployment scenario.

**Table 7.3: Behaviour Aggregate – Example terminal profile mapping**

| Scenario | Behaviour Aggregates |
|---|---|
| Consumer/ SOHO | EF, (AF31), (AF21), (AF11), DF |
| Government/ Corporate | EF, AF31, AF32, AF21, AF22, AF11, DF |
| Backhauling | EF, AF21, (AF31), (AF41), DF |
| Multi-dewlling | EF, AF31, AF21, AF11, DF |
| SCADA | AF11, BE |

The remainder of this clause identifies the main components of the RCST QoS architecture and reviews the relationship between key functions.

## 7.1.1 RCST2 Connectivity Aggregate and Connectivity Channels

Traffic received from the LAN Interface that is to be forwarded by an RCST is divided into one or more Connectivity Aggregates (CA), based on the next hop layer 2 destination to which it is to be forwarded. The term "aggregate" is used generally in the HL user plane to indicate a sequence of HL satellite protocol data units (HLS PDUs). The CA is, hence, the output of a Layer 3 routing or Layer 2 forwarding decision and reflects the interface on which the traffic will be carried over the satellite network (see Figure 7.3). The following examples illustrate the CA concept:

- In a star network, an RCST could use a single CA to forward all traffic towards the GW.
- In a mesh network, an RCST may configure multiple CAs; one could offer connectivity to the GW, and others could offer direct connectivity using mesh connections to other RCSTs.

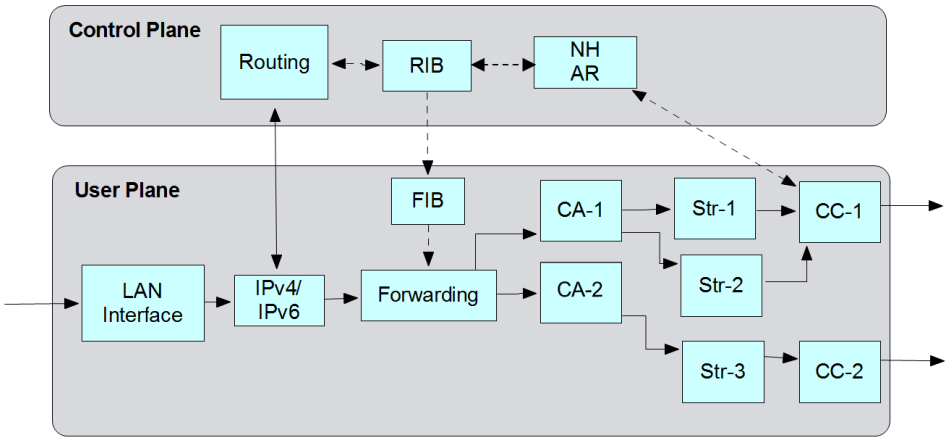


**Figure 7.3: DVB-RCS2 Routing/Forwarding Functions**

Each allocated timeslot is associated with a specific connectivity channel, and hence a single CA. The choice of how many CAs are used depends on how the traffic is to be managed and whether allocated time slots may be used to reach multiple destinations.

Each CA requires a separate instantiation of the QoS framework (i.e. requires an independent set of HLS PDU Queues, a corresponding set of HL Services, set of independently managed QoS and RRM entities, etc.) and supports one Connectivity Channel (CC), which is a physical stream of transmission of bursts. Traffic is routed to a particular CA as a result of a routing decision to a next hop address. Alternatively, a single routing entry could direct traffic to one CA, which later maps traffic to one or more Link Streams. In the second approach, a single set of HL and LL entities may be instantiated.

A CA may utilise multiple Layer 2 streams. Several possibilities exist, for example:

- A single CA does not necessarily imply a single link-layer destination (next hop MAC24 address) or a single Link Stream; this is because DVB-RCS2 allows encapsulating multiple PPDUs with ALPDUs in the same time slot even though these may be destined to different L2 addresses.
- A Link Stream may be used to allow an RCST to identify a mesh destination in a DVB-RCS2 network that supports this connectivity, e.g. when more than one destination is reachable via a CC. The connectivity offered in a mesh system may demand that multiple CAs are used.
- When a CA is used with a multiple access link, it is envisaged that one Link Stream could be configured for each L2 destination.

### 7.1.2 RCST QoS Services

QoS Services are realized in an RCST using a combination of HL and LL Services. CAs are typically subdivided (classified) by a set of TCs that assign the traffic to a specific BA. The TC information may also contain meta-information regarding the traffic specifications for the BA, such as peak-rate, sustainable rate, etc. These values can be used for traffic conditioning (as a DiffServ Policy Enforcement Point) when supported by an RCST. The traffic forming a BA is queued in an HLS PDU queue, which is then mapped to a LL Service Aggregate (SA). After HL and LL processing, the CA will be finally transmitted using a CC. Figure 7.4 presents an example using rectangles to represent functional entities and octagons to represent selection functions. The scheduler (represented by an oval) is an abstract function that determines how HL PDUs are mapped to a SA. Control functions relationships are represented by dashed lines and data flow by solid lines.

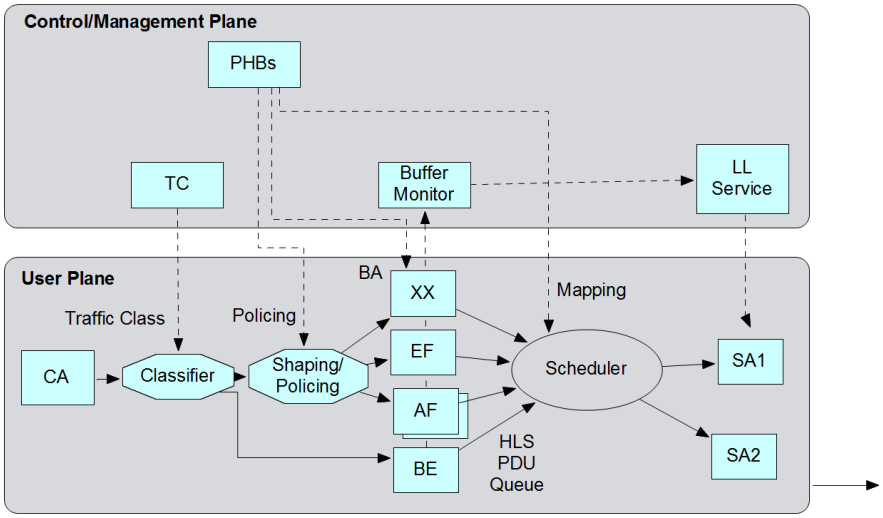


**Figure 7.4: DVB-RCS2 QoS Components**

A more detailed explanation of the operation of the user and control planes is provided in the next clauses.

### 7.1.2.1 User plane QoS

This clause describes QoS processing by an RCST for transmission on the RL. Each RCST has at least one CA that it uses for transmission to the GW. An RCST may also create CAs for other destination within the satellite network (e.g. when supporting mesh connectivity).

Each PDU belonging to a CA is assigned to a single BA, based on a packet classifier that matches the packets to one TC. A TC is implemented as a set of records in the IP classification table (see clause 7.3); each TC matches a set of fields in the IP or L2 header. A classification rule may be as simple as matching only the DSCP or may be more complicated, e.g. involving matching several IP fields with a multi-field (MF)-classifier. A packet classifier may use multiple fields to form a TC:

- Layer 2 information may be used as part of this classification. For instance, a policy may be configured to associate an 802.1pQ PCP value with a specific BA, or a classification rule could be written to assign L2 packets (e.g. LAN control information) to a specific BA based on the Ethertype. This classification may be applied to non-IP traffic.
- At Layer 3, an IP packet may be classified based on the DSCP markings and other IP header information. Together, these fields may be used to select the BA. IP traffic with the Type Of Service field not in line with DiffServ semantics may use the Class of Service (CoS) semantics, rather than those specified by the DiffServ architecture.
- At Layer 4, deep packet inspection may match the port information and other IP payload information to assign the packet to a BA.

Since different levels of classification may result in assignment to a different BA, the RCST needs to specify which fields to trust when there is conflicting information in the TCs.

When a RCST acts as a DiffServ Policy Enforcement Point, the TC may also specify flow properties (e.g. traffic average rate, max burst size, etc.). These properties are used to decide whether the rate of a flow is conformant or non-conformant to traffic specifications. A PDU belonging to a non-conformant flow may be marked (changing the DSCP and/or Explicit Congestion Notification value), and/or re-queued to a different BA, or dropped (discarded). This use implies that a TC may be associated with an additional BA to be used for non-conformant traffic.

Once PDUs are assigned to a BA, they receive the same queuing and IP scheduling treatment (i.e. they are assigned to a single HLS PDU queue). Each BA is characterized by a specific queuing strategy and scheduling method. The traffic forming a BA should be sent using one SA. This set of attributes is collectively referred to as a Higher Layer Service in the control plane.

An SA comprises the PDUs from one or more BAs, and associates these with a priority/precedence. All PDUs assigned to the same SA receive the same treatment by the LL Service. Within the LL, a SA is transported using a Link Stream (LS) that carries a sequence of L2 packets. For example, a LS may be associated with Payload-adapted PDUs (PPDUs) of a LL logical flow. Packets are finally multiplexed into bursts or Transmission Streams (TX Streams) for transmission over the air interface.

The precedence of a LL service is used to inform scheduling decisions when a transmission opportunity is made available to an RCST. When more than one SA is defined L2 pre-emption may be used. This allows the QoS system to suspend transmission of a partially-transmitted PDU (Link Stream packet) from a lower priority SA at the end of a transmission burst, and initiate (pre-empt) the next transmission burst with a PDU from a higher priority SA. The transmission of the lower priority SA is resumed in a later timeslot. This method can be used to upperbound the jitter experienced by higher priority SA traffic.

#### 7.1.2.2 Control plane QoS

The HL Service is defined as a per-hop treatment of Layer 3 PDUs characterized by a PHB. This is a management construct that includes the policy needed to instantiate the PHB and relate this to an HLS PDU queue. Each HL Service corresponds to one BA. This defines the parameters that are needed to support PHB-specific operations, including queuing and scheduling, and the SA to be used.

A single PHB may be instantiated to form multiple instances of the HL Service. Conversely, a single HL Service may be used to support multiple PHBs using a single BA. However, since the HL Service can not differentiate the treatment of PDUs within the BA, a set of HL Services need to be instantiated to realize DiffServ QoS. A consistent QoS treatment across HL and LL is guaranteed by defining one LL Service (SA) for one or more HL Services. Each LL Service is created by the LL Service Descriptor in DVB-RCS2 Lower Layer Specification.

The LL Service provides an interface to access the satellite resources. The configuration of the LL services associated to a SA and its corresponding Link Streams determines:

- the allowed mapping between Link Streams and Request Classes (RCs),
- the allowed mapping between Link Streams and dedicated-access allocation channels,
- the allowed mapping between Link Streams and random-access allocation channels.

An LL Service specifies the types of Allocation Channels (AC) that may be used for each SA. The AC identifies a portion of the RL capacity that is available for use by one or more LSs:

- The Dedicated Access AC (DA-AC) receives allocation by means of explicit demand/assignment methods or free capacity assignment (FCA).
- The Random Access AC (RA-AC) represents a portion of the return link spectrum that is offered for random access for multiple terminals. NCC may use a load-control algorithm to control the level of contention on the RA channel.

An LL Service may allow an SA to use one or more (DA-ACs) LL Service or one or more RA-ACs. In addition, it may optionally be mapped to other AC, with each AC mapped to a connectivity channel. The LL Service can also inhibit access to the DA-AC or RA-AC, for instance because an RCST does not support these LL Services.

An RCST uses an AC to select a specific RC or for QoS differentiation:

- The LL Service contains a reference to the AC and the RC. This defines implicitly an association between the AC and the RC. When the RCST generates a Capacity Request (CR), it inserts the RC identifier (RC_index) into the CR message to communicate this value to the NCC.
- An AC is also used to differentiate connectivity channels, when multiple capacity categories are used.

The NCC is responsible for distribution of DA-ACs and RA-ACs. The NCC can address a specific DA-AC by inserting an Assignment_ID in the TBTP2.

- Each DA service is the control plane correspondence for the user plane DA-AC. This is specified in the RCST configuration. As seen from an RCST, there is a one-to-one correspondence between a DA service and a DA allocation channel.

- The RA service is the control plane correspondence for the user plane RA-AC. This is defined by the resources provided to the associated RA-AC (as controlled by the NCC), the RA Load Control parameters associated with the RA-AC and the current loading of the RA-AC by the RCST.

An RCST assumes that a DA Service will be assigned capacity by the NCC as specified for the nominal RC associated with the DA Service. The DA Service specification can then be inferred from the configuration of this RC.

#### 7.1.2.3 Management plane QoS

The HL Service mapping contains a number of managed QoS parameters that characterize the HL Service, such as: MinRate, MaxRate, MaxIngressBurst, MinIngressBurst, MaxDelay, MaxLatency and LinkRetransmissionAllowed. Also, it contains information relating to the queue behaviour, such as SchedulingType.

Table 7.3 in [i.1] provides the minimum set of HL Service parameters. Additional parameters could be added by the implementor, if necessary, to better specify the expected QoS behaviour of the user PDU within the RCST.

In addition, in Table 7.4, some examples are provided for queue configuration for some BA / HL Service.

**Table 7.4: Example HL Service Configuration**

| HL Service (PHB) | Dropping Mechanism | Scheduling Type | Queue size |
|---|---|---|---|
| EF | Tail Drop | FIFO | MaxIngressBurst*MaxDelay |
| AFxx | Random Early Detection | WFQ | MaxIngressBurst*MaxDelay |
| BE | Tail Drop | FIFO | MaxIngressBurst*MaxDelay |

Traffic is classified by matching against a set of TCs, an example is shown in Table 7.5.

**Table 7.5: Example TC Service Configuration**

| MF Classifier | HL Service | Metering System |
|---|---|---|
| EF DSCP (46) | EF | Disable |
| AFxx DSCP | AFyy | Single Rate Three Colour Marker;<br>Two Rate Three Colour Marker |
| BE DSCP (0) | BE | Disable |

## 7.2 QoS organization configuration

The organization of BAs and Link Streams and the distinction of LL and HL Services provide a number of ways to configure the QoS support in a RCST.

Two distinct mappings can be identified for a star system:

- A mapping may regard all allocated timeslots as belonging to a single SA. This is the simplest method. It queues packets by BA within the HL, and requests capacity using one or more RCs. Although different policies may be used to request capacity for different TCs, all allocated timeslots are used as one service by the scheduler, which optimizes use according to the assigned LL Service.

- A mapping may provide a strict separation between a set of LL Services. This queues packets by BA within the HL, and also may request capacity using more than one RC. Different policies may be used to request capacity for different TCs. The allocated timeslots are differentiated at the scheduler by LL Service, which seeks to assign the traffic to the allocations made in response to each RC. A policy may be used to reassign unused timeslots to other traffic.

Figure 7.5 illustrates two organizations of the QoS system in a RCST. DVB-RCS2 does not specify a particular method for scheduling. In this example, a scheduler is assumed to be triggered by allocation of a timeslot.

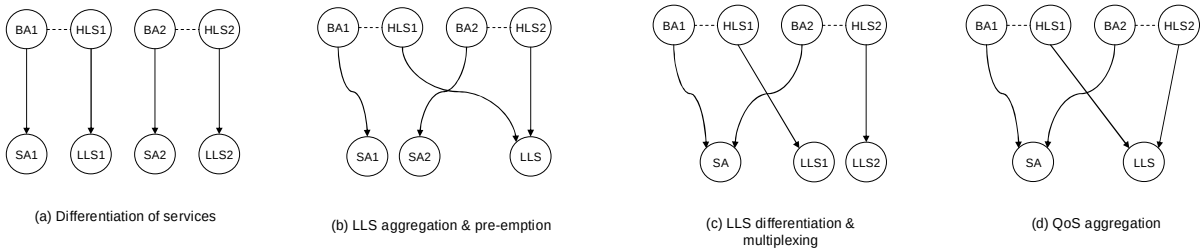


**Figure 7.5: Example QoS mappings**

An example of the first organization is shown in Figure 7.5(a). Two BAs are each mapped to a SA. The HL Service associated with each BA is also mapped to a distinct LL Service. This organization can be used to support two distinct services that operate independently. Each LL Service could be independently billed, policed, and cannot be adversely impacted by traffic assigned to another LL Service. Independent L2 allocation methods are used to request transmission resource for each BA, and the NCC using a corresponding assignment_id identifies the allocated timeslots. The RCST scheduler will use the assignment_id to schedule SA traffic.

A different organization is achieved when the BAs are mapped to separate SAs, but their respective HL Service is multiplexed to a common LL Service (Figure 7.5(b)). This organization allows the scheduler to use the optimum policy to schedule the use of the allocated capacity (according to the parameters set in the HL Service for each BA. Since different Link Streams are used at L2, the L2 scheduler is responsible for determining the order of transmission of packets and may support pre-emption of lower priority SAs.

Other organizations are also possible using the RCST QoS architecture.

### 7.2.1 Scheduling in RCST

DVB-RCS2 intentionally does not specify the semantics of scheduling in RCST. This leaves implementers the flexibility to perform scheduling decisions either in the lower layers, the higher layers, or a combination of the two.

In the following example, it is assumed that a scheduler is used that is triggered by a transmission opportunity on either a RA or DA channel. This is most easily envisaged in the DA case, where the task of the scheduler may be to determine whether there is any data that should be sent in the transmission opportunity, and if so what data should be sent. The data to be sent comprises:

- Partially sent HLS PDUs pending completion of the final fragment(s)
- HLS PDUs queued at the higher layer

If two BAs are mapped to the same SA, scheduling decisions are made before packets are encapsulated at L2 (Figure 7.5 (c) and (d)). However, distinct allocation methods can still be used for the two BAs. This organization allows allocation of resources to different streams of traffic (e.g. VoIP, web traffic, interactive services, etc.) within the same allocation pool. In this case, L2 pre-emption may not be required.

### 7.2.2 Example use of RCST QoS system model

Figure 7.6 illustrates the relationship between modules the higher layer QoS functions and the lower layers QoS functions. The diagram is intended to be informative and does not mandate any particular internal structure of an RCST. Solid lines represent the flow of PDUs and other data through the system, whereas dashed lines are used to denote control relationships. Simple functions or objects are represented by boxes, selector mechanisms by hexagons, and complex objects by pentagons.

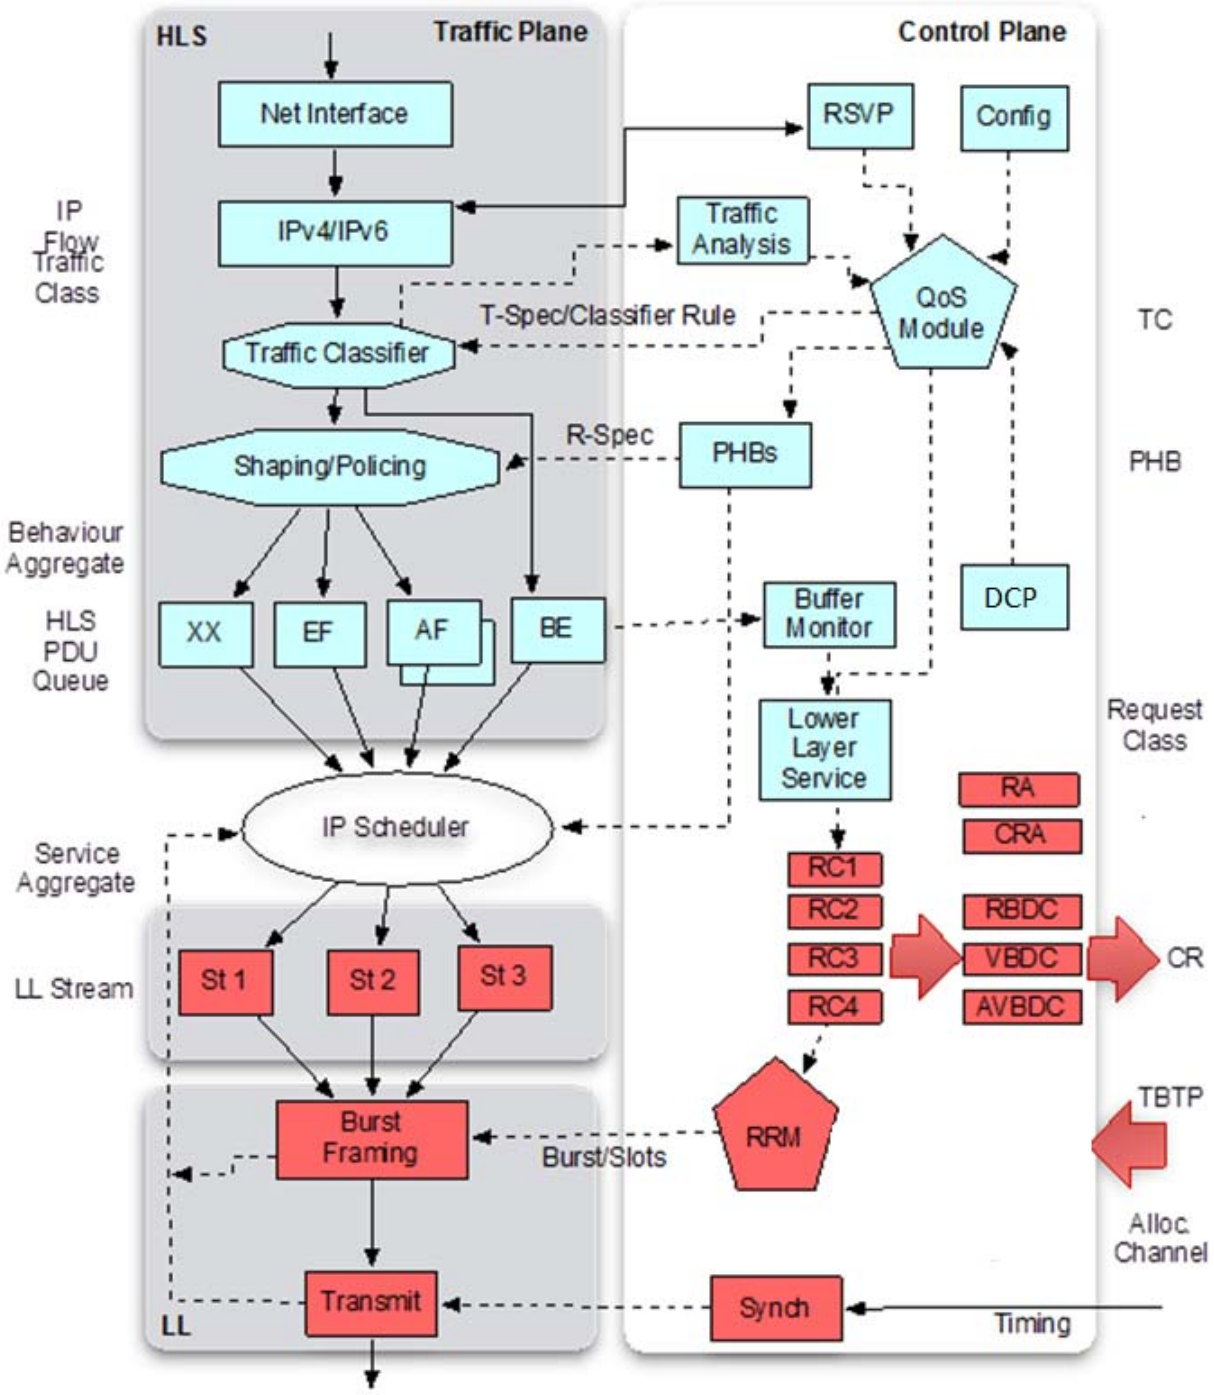


**Figure 7.6: Logical HLS QoS Processing**

In the diagram, the data paths are represented by dashed lines and control paths by dashed lines. Traffic arriving at the LAN network interface of an RCST has been divided into several Traffic Classes (TCs). These classes are mapped to 5 per-hop behaviours (PHBs). These traffic classes may for instance reflect a best effort Diffserv Code Point (TC1), and unknown service category (TC2) – in this case mapped to the Best Effort (BE) PHB, an Assured Forwarding codepoint mapped to one of the two AF PHBs, and an Expedited Forwarding class mapped to the EF PHB. The final traffic class maps to be a special-purpose class, the XX PHB. Each HLS PDU queue (behaviour aggregate) is in turn mapped to a Link Stream (service aggregate) for transmission (ST1-ST3). The Radio Resource Management (RRM) object is responsible for requesting capacity from the NCC.

The outputs of the HLS PDU Queues hold the data to be sent over the lower layer service. This implies the action of an IP scheduler (represented by a white oval). This may be understood to be activated each transmission opportunity (notified by the TBTP2) to select the PDUs that are segmented into the stream. The selection is based on the PHBs (which indicate the lower service), and link-layer information. This ensures that PDUs or segmented PDUs are sent using the corresponding allocation channel. When required, PDUs pass through a segmentation function, so that any unsent data is postponed to a later scheduling opportunity. Each segment is then encapsulated into one of the configured streams (ST1-ST3 in the diagram) and is then placed in the burst for transmission. The scheduler could use a strict priority scheduler or a weighted priority scheduler, but is not specified in the present document. Since in this example there are three Link Streams, ST1 can preempt ST2 or ST3.

## 7.3 QoS configuration management

The RCST basic QoS configuration may be provided in the new configuration file download or via TIMu NLID messages. It may also be provided directly by the installer (e.g. by manually configuration or via local configuration file download).

The first time that the RCST enters the system, it is recommended to always verify the full HLS configured data. The exact value of the QoS HL parameters will depend on the RCST profile and should follow the recommended values provided in previous clauses. The NMC side could have different QoS templates depending on the RCST profile in the system. This information should be part of the RCST commissioning together with other HL parameters.

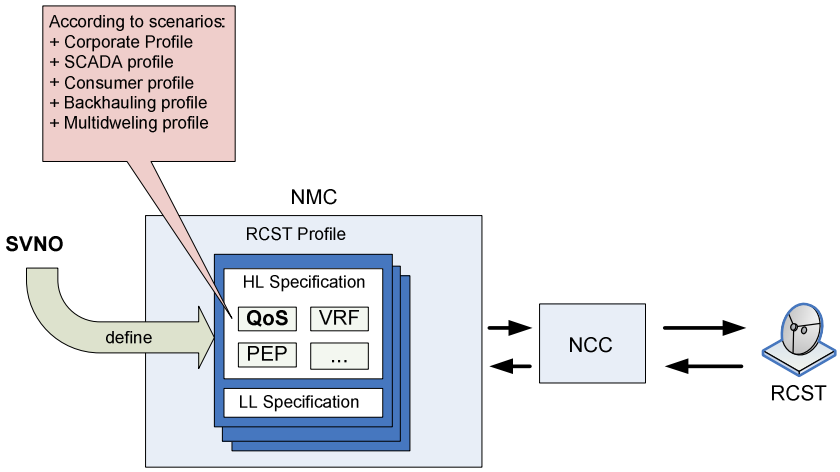


**Figure 7.7: QoS configuration management**

The HLS QoS configuration should, at least, contain the following information:

- IP Classification table: This defines a TC in the form of a table that maps each PDU to a specific BA,. If there is no entry in this table, then there is no way to classify the traffic, and the RCST by default may drop the user traffic. The exact number of entries in IP classification table will be system dependent. However, at least one default entry should be provided (e.g. to a best-effort BA).

- HL Service mapping table: This maps the HL services to one LL service for consistent QoS treatment. This table should contain at least one HL service per LL service provided during logon messages.

The setting of NMC QoS parameters per RCST is given by the SVNO. The SVNO is responsible for the traffic functions, IP routing, and QoS. Therefore the SVNO should provide the RCST QoS HLS templates per profile in the NMC. The NMC should contain different RCST profiles depending on the terminal profile type of RCST that the SVNO works with. It may be possible to have more than one template per profile, which will be a system implementation decision.

## 7.4 QoS management and control in regenerative mesh networks

### 7.4.1 DVB-RCS2 Logon with regenerative mesh support

Dynamic connectivity support, connection control protocol support and version, and transparent mesh capabilities are included in the LL capabilities and HL capabilities groups of the logon element types sent by the RCST.

The Logon TIMu includes the Logon Response Descriptor, which may provide one dedicated access allocation channel (DA-AC) for control and management, providing the resources for the mesh signalling connection. This is done by associating an Assignment ID to signalling. More information to establish the signalling connection may be provided in the NLID descriptor.

The Logon TIMu also includes the Lower Layer Service Descriptor, where the allocation channel applying to each lower layer service, and its corresponding RC (Request Class), are indicated. For multi-beam mesh systems, a different allocation channel per physical destination may be needed. This is due to the possibility that the physical resources associated to the different destination are disjoint. Such is the situation when the satellite switching is performed at Layer 1. In this case, different Assignment IDs can be assigned per destination beam supported by the RCST. The RC associated to each AC can be independently configured, according to the LL Services configuration.

Configuration of default values for the HL QoS tables (IP Traffic Classification and HL Service tables) might take place by reception of NLID descriptor (in the form of SNMP set commands) included in the Logon TIMu.

### 7.4.2 HLS Maintenance

The SVNO may (re)configure or add new entries in the IP Traffic Classification and HL Service tables of the RCST, using SNMP protocol or other IP-based method (e.g. configuration file download).

Certain parameters of the HL Service table entries will not be modifiable by the SVNO, as the LL service associated to an HL Service.

The default entry of the IP Traffic Classification and HL Service tables should not be modified or deleted by the SVNO.

### 7.4.3 QoS Configuration for regenerative mesh systems

The RCST QoS data structures used for regenerative mesh systems are:

1) Traffic classification and services tables ():
   - IP Classification Table (defined in [i.1]).
   - HL Service Mapping Table (defined in [i.1]).
   - LL Service Table, created after reception of the LL Service descriptor at logon.
   - RC Table, created after reception of the LL Service descriptor at logon.
2) DCP specific tables:
   - Active Connection Table.
   - Connections timeout value. This value may also be established by using DCP.

The Active Connection Table (see Table 7.6) lists the active DCP connections and needs to include the following parameters:

**Table 7.6: DCP Active Connections Table**

| Parameter | Description |
|---|---|
| ActiveCnxIndex | The index in the DCP active connections table. |
| ActiveCnxRefId | A reference for the DCP connection. |
| ActiveCnxMACSrcAddr | MAC24 address of the originating interface in the source RCST. |
| ActiveCnxMACDestAddr | MAC24 address of the destination RCST reception interface, or multicast MAC24 for unidirectional multicast connections. |
| ActiveCnxType | Unidirectional (unicast or multicast) or bidirectional. |
| ActiveCnxService | QoS Service for the connection (used for transmission). |
| ActiveCnxAssignmentId | Reference to map TBTP2 resources assigned to the connection for this RCST. |
| ActiveCnxOtherAssignmentId | Reference to map TBTP2 resources assigned to the connection for the peer RCST. |
| ActiveCnxIPv4SrcAddress | In the source RCST, IPv4 address of the LAN interface that received the IP traffic packet triggering the connection. |
| ActiveCnxIPv4DstAddress | IPv4 address of the peer RCST for this connection, taken from the Connection Establishment Request message. |
| ActiveCnxIPv6SrcAddress | In the source RCST, IPv6 address of the LAN interface that received the IP traffic packet triggering the connection. |
| ActiveCnxIPv6DstAddress | IPv6 address of the peer RCST for this connection, taken from the Connection Establishment Request message. |

### 7.4.4 QoS MIB Objects for regenerative mesh

The QoS parameters, that determine the QoS profile and Allocation Channel (AC) corresponding to one mesh connection, can be extracted from the HLS tables (IP Classification and HL Service MIB tables), included in the *dvbRcs2QoSConfiguration* MIB group [i.1] and the LL Service Descriptor (LL Service and RC MIB tables).

Table 7.2 in [i.1] includes the defined IP TCs. This table links to the HL services table through parameter *IPClassHLSAssociation*. The TC entry may apply to only one satellite SVN, if its associated HL service is defined for a specific interface. The RCST may discard VLAN frames if their user priority does not match the value specified by *IPClassVlanPri* parameter. Also, the RCST can drop IP packets based on IP header values through parameter *IPClassAction*. In this way it can be avoided that a user packet triggers a DCP connection request, the table entry is then acting as a reverse firewall.

Table 7.3 in [i.1] characterizes HL services and links them with LL services. Each entry of the HL Service table is associated to a LAN interface number of the RCST, which will map to the *interfaces* MIB group, being applicable only to one SVN (linked to a VLAN_ID). This means that HL Services should be replicated (when necessary) for each SVN supported by the RCST. The recommendation is that the *IPClassVLANID* field in the *IPClassTable* is left empty when the entry of the table applies to all the SVNs supported by the RCST. The VLAN mapping table (Table 8.14) maps user VLAN_IDs and satellite SVNs.

Table 7.7, constructed from the LL Service Descriptor in the Logon Response, maps LL services with RCs and ACs. A reference for the LL service that the RCST intends to use for the mesh link will be included in the DCP establishment request message. This LL service maps to one RC class.

**Table 7.7: LL Service parameters**

| Element | Description |
|---|---|
| LLserviceIndex | Index of LL service Table. |
| LLserviceRCIndex | A 4 bit field indicating the nominal request class for the associated Link Service. |
| LLserviceDAACIndex | A 4 bit field indicating the nominal dedicated access allocation channel associated with the Link Stream. The Assignment ID associated to the request class has an offset to the Assignment ID Base equal to the nominal_da_ac_index. |
| LlserviceCD_RCmap | A 16 bit field indicating the allowance to conditionally map resource demand for the associated Link Stream into capacity requests for other RCs, with bit 0 referring to rc_index=0, bit 1 referring to rc_index=1 and so on. |
| LLserviceCS_DAACmap | A 16 bit field indicating the allowance to conditionally map traffic from the Link Stream into the different dedicated assignment allocation channels, indicated by a flag for each DA-AC, with bit 0 referring to da_ac_index=0, bit 1 referring to da_ac_index=1 and so on. |

Table 7.8, derived from the LL Service Descriptor, defines the RCs in the system that are usable by the RCST. These RCs are used by mesh links CCs according to the QoS service selected by the RCST originating the request.

**Table 7.8: RC table parameters**

| Element | Description |
|---|---|
| RCindex | The RCST by default maps its default request class to rc_index 0 |
| RCcontantAssignment | Flag to indicate if constant non-solicited assignment is provided for the RC<br>Values: Non-solicited(0), Solicited(1) |
| RCvolume_allowed | Flag to indicate if A/VBDC requests are allowed for the rc_index<br>Values: NotAllowed(0), Allowed(1) |
| RCrbdc_allowed | Flag to indicate if RBDC requests are allowed for the rc_index in kilo bits per second<br>Values: NotAllowed(0), Allowed(1) |
| RCmax_service_rate | Field that indicates the maximum service rate for the rc_index. The maximum allowed RBDC equals this level substracted by the CRA in kilo bits per second |
| RCmin_service_rate | Field that indicates the minimum rate that can be expected assigned when actively requesting any dynamic capacity for the rc_index |
| RCconstant_service_rate | 16-bit field indicating the admitted CRA level associated with the request class in kilo bits per second |
| RCmax_backlog | 8-bit field indicating the max volume request backlog that the NCC will accept to hold for the rc_index in kilo bits per second |

# 8 Satellite Virtual Networks and VLANs

## 8.1 Mapping of SVN tags to lower layer fields

This clause provides guidelines on the SVN tag mapping to lower layer fields. The concept of SVNs is explained in [i.1]. At layer 2, each logical RCST network interface towards the satellite system has a unique 24-bit MAC24 label that consists of an SVN-number or SVN-prefix and an SVN interface ID. The boundary between the two is variable and configured via L2 signalling and may be different for difference SVNs running in the same system. When using different SVN prefix lengths, care should be taken to assign addresses so that the SVN numbers are non-ambiguous. The allowed range of the prefix length is 1 to 16 bits, so that an RCST should support at least two SVNs – SVN 0 for management and one or more SVNs for user traffic.

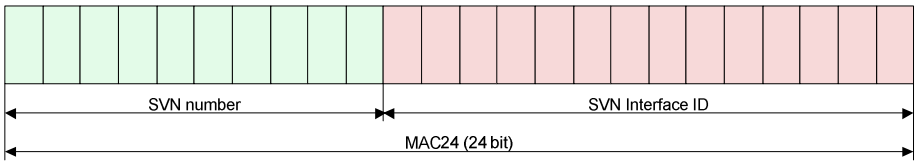


**Figure 8.1: Format of MAC24**

The length of the SVN number (the boundary) is signalled in the L2 Logon Response Descriptor. This descriptor can configure up to 15 RCST addresses, the SVN number length (MAC24_prefix_size) is set for each of them independently and care should be taken that no overlapping addresses are created. The prefix length field svn_prefix_size has a size of 5 bits, but its contents are restricted to the range [1..16].

The actual value should be decided based on the maximum number of RCSTs to be supported in the given SVN. If, for example, up to 2 048 RCSTs are to be supported, the interface ID size should be at least 11 bits and the svn_prefix_size should be less than or equal to 13.

The Logon Response Descriptor also configures the assigned MAC24 label in the unicast_mac24 field, and the default SVN number in default_svn_number.

When encapsulating a higher layer PDU, the RCST should compare the SVN number of the packet to the default_svn_number. If these match, the RCST should use a 0-byte packet label for encapsulation (i.e. label type 2). If they do not match, the RCST should take most significant byte of the MAC24 and place this into the ALPDU label with label type 0.

### 8.1.1 MAC24 address assignment to terminals

The MAC24 addresses assigned to a terminal should provide non-ambiguous mapping to and from SVN/interface-ID.

**00000000 00000000** 00010001/16
**00000001** 00000000 01010101/8
**00000010 0001**0001 00000011/12
**1111**0001 00000001 00000011/4

**Figure 8.2: MAC24 assignment example**

The upper 8 bits of all MAC24 addresses assigned to a terminal explicitly or implicitly (multicast addresses) should be unique as in Figure 8.2. The number after the slash is the SVN-prefix length and the bold digits are the SVN-prefix. When sending a higher layer PDU to using any of the four addresses the ALPDU label (the upper 8 bits) will be different in each case. This means provided the hub receiver with the means to decide to which of the four SVNs the packet belongs.

### 8.1.2 GSE transmitter processing

When sending a higher layer PDU, the hub puts the complete MAC24 address into a 3-byte label of the GSE-encapsulated PDU (label type 1). For unicast packets, this is the MAC24 address assigned to the corresponding SVN of the destination terminal, for multicast packets this is a multicast MAC24 address created via the selected multicast addressing method. Layer 2 M&C messages are labelled with either the MAC48 of the terminal address, if they are unicast (label type 0), or with no label, if they are broadcast (label type 2).

### 8.1.3 GSE receiver processing

After reassembly and decapsulation, the receiver filters the packets as follows:

- If there is a 6-byte label and this matches a MAC48 assigned to the terminal and the protocol type is L2 M&C signalling, the packet is accepted and forwarded to the signalling handling function.
- If there is no label and the protocol type is L2 M&C signalling, the packet is accepted and forwarded to the signalling handling.
- If there is a three byte label and it matches one of the MAC24 addresses assigned to the satellite-side interfaces of the terminal, the packet is accepted. The SVN-number is extracted from the address by appropriate masking, extended to 16-bit by adding zero bits at the LSB (if necessary) and forwarded to the higher layer functions together with the corresponding SVN.
- If the packet is associated with SVN 0, then the packet is passed to the HLS management function.
- Otherwise the packet is dropped.

Note that the filtering is conceptually performed after reassembly and decapsulation. It is possible to do the filtering before these steps by applying a cache technique to map fragment ids to labels. This is expected to have the same forwarding behaviour.

An example of the forwarding to higher layers follows. Given that the MAC24 assignments in Figure 8.2 the processing shown in Table 8.1 takes place.

**Table 8.1: RLE receiver processing of labels**

| Label on GSE packet | SVN | SVN passed to HLS |
|---|---|---|
| 00000000 00000000 00010001 | 00000000 00000000 | 00000000 00000000 |
| 00000001 00000000 01010101 | 00000001 | 00000001 00000000 |
| 00000010 00010001 00000011 | 00000010 0001 | 00000010 00010000 |
| 11110001 00000001 00000011 | 1111 | 11110000 00000000 |

### 8.1.4 RLE transmitter processing

The following processing takes place on the terminal sender side:

- If the higher layer packet is a L2 M&C signalling packet, label type 3 is used and the packet is encapsulated with suppressed protocol type.
- If the packet is another higher layer PDU, its SVN number is compared to the SVN numbers of the MAC24 addresses assigned to the terminal (the SVN prefixes padded to 16 bit with zero bytes at the LSB).
  - If there is no match between the SVN number of the higher layer PDU and the SVN numbers of the MAC24 addresses assigned to the terminal, the packet is dropped.
  - The matched SVN is compared to the default SVN signalled by the hub. If these match the PDU is encapsulated with label type 2 and no label.
  - If the SVN does not match the default SVN, the PDU is encapsulated using label type 0 and the most significant 8 bits of the MAC24 address are placed in the 1 byte ALPDU label.

### 8.1.5 RLE receiver processing

The following processing takes place at the hub receiver after reassembly and decapsulation:

- If a packet arrives with label type 3 and a protocol type of L2 M&C signalling, it is passed to the signalling entity in the hub with the terminal id attached.
- If the packet arrives with no label, the default SVN (in extended form 16-bit) is attached to the PDU and the PDU together with the SVN number are passed to the higher layers.
- If the packet arrived with one byte label, the MAC24 address assigned to the terminal is searched which has the label byte as its most significant byte.
  - If there is no match, the PDU is dropped. This event should be logged as it is a symptom of misconfiguration.
  - If there is a match, the PDU and the SVN number (in extended form of 16-bit) are passed to the higher layers.

## 8.2 Recommendations for VLAN support and Satellite Virtual Networks

In this clause, usage examples of VLAN and SVN are described.

The 3B VLAN (IEEE 802.1pQ [i.72]) tag contains three fields: a Priority Code Point, PCP (3 bits), a Canonical Format Indicator, CFI (1bit) and a VLAN identifier, VLAN_ID (12 bits).

Ethernet frames that are received at a VLAN-enabled LAN interface at the GW or RCST may contain an IEEE 802.1pQ [i.72] tag or may be untagged. Untagged frames received at a VLAN-enabled LAN interface should be associated with a default VLAN_ID.

In the rest of this clause, a "LAN interface" always refers to a VLAN-enabled LAN interface.

VLAN and SVN support in DVB-RCS2 systems may be realized by two methods:

a) In this method, each VLAN tag, which is associated with a LAN interface, is mapped to a specific MAC24 address. When an Ethernet frame is received with a given VLAN tag, the tag is removed, and the frame is mapped to a specific MAC24 address. This requires an explicit configuration of the VLAN_ID used at the ingress and egress LAN interfaces. The egress interface may use an untagged format or add an 802.1pQ [i.72] tag. If 802.1pQ tags are used at the egress LAN interface the VLAN PCP should also configured for this interface, since there is no PCP value in an untagged Ethernet frame In this case, the Ethernet PCP may be mapped from the IP DSCP. Thus configuration is needed to assign the traffic to VLANs and set the PCP codepoint
(e.g. static configuration via the management plane, or a dynamic method using the control plane).

b) Forwarding of an IP packet with an IEEE 802.1pQ [i.72] tag from/to an RCST that operates as a router. In this method, the RCST operates as a router, and it forwards IP packets with their IEEE 802.1pQ tags. In this case, a SVN interface may be configured so that RLE/GSE headers carry the value of the VLAN-ID as a tag. Multiple VLANs are implicitly identified by the encapsulation tag value and may use the same MAC24. This mode still requires configuring the VLAN-PCP at the egress LAN interface for the 802.1pQ tag.

In method (a), the RCST LAN interface should be configured with a corresponding MAC24 address for each VLAN-ID that is supported. The Gateway should also be configured with a corresponding VLAN-ID or a separate interface for each VLAN. The RCST and Gateway VLAN-IDs for the same VLAN may be different. The possibility of maintaining different VLAN_IDs at RCST and GW LAN interfaces for the same VLAN enables the Gateway operator to separate traffic that is carried using the same VLAN-ID at two RCSTs, but where this value is intended to identify two independent networks. Frames with a VLAN-ID value that has not been configured, should be dropped at the ingress interface.

In method (b), the RCST LAN interface is configured with the set of VLAN-ID it supports. In this method, a single MAC24 address may be used to forward more than one VLAN-IDs over the satellite interface. In this method, the Gateway operator should either enforce a policy on the use of VLAN-IDs at RCSTs (ensuring that each VLAN-ID identifies only one satellite VLAN), or use a separate interface for each separately managed set of VLANs. This follows normal practice for Ethernet-based VLANs. A VRF group may also be used for this purpose, since each VRF group is presented on a separate interface at the Gateway.

If VLAN support is realized through method b), GSE/RLE encapsulation, clause 8.6.21 "VLAN configuration group" of the RCS2 specification [i.1] indicates that the following MIB table entries are needed:

- A management parameter describing if an RCST is capable of supporting method-(b).
- A management parameter for the NCC to control a specific LAN interface of the RCST.

If the RCST enables VLAN method (b), the following MIB table entries need to be configured, for the LAN interface:

- A set of allowed VLAN_IDs. The frames corresponding to the configured VLAN-IDs are mapped to the corresponding SVN. A default interface may be configured to forward frames received with a VLAN_ID not specified in this set of allowed VLAN-IDs.
- A maximum Priority Code Point (PCP) value. A higher value indicated in the VLAN tag will be truncated to this value. This rule may be used to enforce operator-controlled use of the PCP values, for example to reserve the highest values for specific groups of customers or specific applications.

Method (a) is the method used in a routed IP network. Method (b) extends the concept of a VLAN across the satellite network, forming a hybrid of a L2/L3 network that allows coordination of the VLAN-ID values used over the networks connected via the RCSTs.

It is common for ISPs to offer a single LAN interface to the subscriber LAN. VLAN services to individual subscribers are not common. It is expected that multi-VLAN support at the LAN interface of the ST will be attractive where isolation between different LAN users is needed. Thus, the VLANs may be terminated at the RCST using IP routing over the satellite air interface, or extended across the satellite network using either method (a) or (b) to assign the VLAN_ID. From the point of view of the connectivity, transparent and regenerative architectures may be enabled.

Examples of VLAN usage are described in the following clauses.

### 8.2.1 Consumer/SOHO scenario

Consumer networking equipments usually do not support VLANs. It is common for ISPs to offer a single LAN interface to the subscriber LAN. VLAN services to individual subscribers are not common and are not required in this scenario. A RCST will likely be a part of only one traffic SVN.

### 8.2.2 Corporate/Institutional (including Military) scenario

Corporate and Governmental networks frequently use VLANs to segregate traffic between user communities and often employ IP routing to connect VLAN-enabled LANs. VLAN support is therefore expected for this scenario, where the RCST may be part of one or more traffic SVNs. A range of configuration examples is given below.

#### 8.2.2.1 Configuration example 1

Figure 8.3 shows example 1, the case of two RCSTs that belong to the same traffic SVN (SVN-1) with each RCST supporting one VLAN (VLAN-1). The LAN interfaces at each of the RCSTs should be configured to associate the traffic with the same VLAN_ID. The SVN for management (SVN-0) is also shown, as well as its respective VRF group (VRF-0).

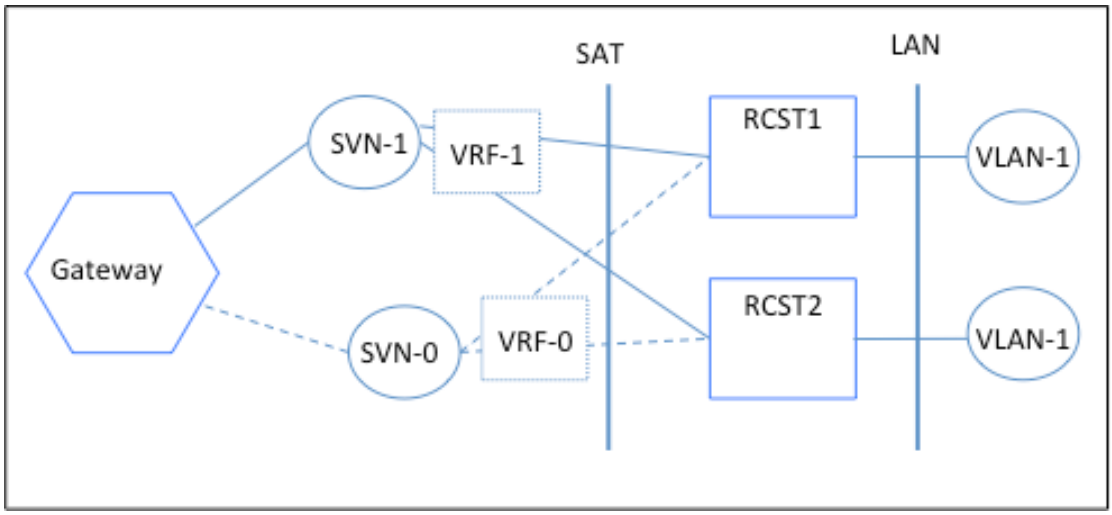


**Figure 8.3: Example 1 - Two RCSTs in one SVN; each RCST supports one VLAN**

An example of a lookup table configured by the INAP/SVNO, when method a) is supported, for the topology of Figure 8.3 is shown in Table 8.2. An SVN_MASK length of 8 bits (e.g. 255 SVNs can be supported by the OVN) is considered, although other sizes are also applicable.

In the topology of Figure 8.3, Ethernet frames (with no tag) from VLAN-1 of RCST-1 will be forwarded through the LAN interface 1, associated momentarily in the RCST with VLAN_ID 1 and, will then be assigned an MAC24 label of 0x1000A1 corresponding to SVN-1. Then, in this example, the LAN supported by RCST1 does not use VLAN tags (i.e. frames with IEEE 802.1pQ [i.72] tags). Tagged packets arriving to RCST1 will be dropped.

In contrast, the LAN at RCST2 has been configured with VLAN support. Tagged frames with a VLAN_ID of 1 that are received by RCST2 will be assigned a MAC24 label of 0x1000B1 corresponding to SVN-1, and will be forward via the satellite (with prior removal of their 802.1pQ tag). Untagged packets arriving at the LAN interface of RCST2 will be dropped.

All the traffic is carried in one VRF group, and hence could be presented at the Gateway using an interface with 802.1pQ to identify each VLAN. Packets received by the Gateway that correspond with SVN-1 will be mapped to the configured VLAN for the SVN, before being forwarded to the Gateway LAN interface. In this case, they are mapped to VLAN-1 (any other VLAN_ID may be mapped, including an untagged value). In this case, SVN-0 is mapped to a separate interface port for management data, because it belongs to a separate VRF group. The use of a separate VRF Group means that the addresses and any created VLANs in a single VRF Group are completely independent of any in other VRF Groups. (Hence, VLAN_ID 1 in VRF-0 (the management VRF Group) is entirely independent of VLAN_ID 1 in VRF-1 (the first traffic VRF Group).

SVN-1 is mapped to VLAN_ID 1 for traffic. RCST1 uses a non-tagged format at its LAN Interface, whereas RCST chooses to encapsulate the traffic sent on the LAN Interface using VLAN-tagging. In both cases the RCSTs will associate the traffic with VLAN_ID 1.

Thus, this is an example of SVN/VLAN support using method (a), i.e. no implicit coordination of VLAN_IDs between the LAN interfaces of the STs and the Gateway.

**Table 8.2: Example of VLAN mapping to support SVN/VLAN using method a) for Figure 8.3 (default = with no tag on the Ethernet LAN interface; tagged = with a 802.1pQ tag)**

| | MAC24 | VLAN_ID | Interface |
|---|---|---|---|
| Gateway | SVN-0: 0x000001/8<br>SVN-1: 0x100081/8 | VRF-0/VLAN-1<br>VRF-1/VLAN-1 | LAN 0 (Mgmt)<br>LAN 1, VLAN-1 |
| RCST1 | SVN-0: 0x000002/8<br>SVN-1: 0x1000A1/8 | VRF-0/VLAN-1<br>VRF-1/VLAN-1 | Internal (Mgmt)<br>LAN 1,default |
| RCST2 | SVN-0: 0x000003/8<br>SVN-1: 0x1000B1/8 | VRF-0/VLAN-1<br>VRF-1/VLAN-1 | Internal (Mgmt)<br>LAN 1, VLAN-1 |

Table 8.3 shows an example configuration with support of VLANs using method b). Untagged frames, received at RCST-1, will be tagged with a default VLAN_ID of VLAN-1 and the MAC24 of 0x1000A1. If their PCP field is lower or equal to 5, there will be no change to this value; otherwise the RCST will modify the PCP value reducing it to 5. For RCST2, tagged frames will be also tagged with the VLAN_ID of VLAN-1 but will use a MAC24 of 0x1000B1. The maximum PCP value for the traffic of this may be different, and in this case is 4.

Note that a single SVN could be used to support multiple VLANs.

**Table 8.3: Example of VLAN mapping to support SVN/VLAN using method-b) (default = without tag on Ethernet LAN interface; tagged = with a 802.1pQ tag)**

| | MAC24 | VLAN_ID | Interface | PCP |
|---|---|---|---|---|
| Gateway | SVN-0: 0x000001/8<br>SVN-1: 0x100081/8, tagged | VRF-0/VLAN-1<br>VRF-1/VLAN-1 | LAN 0 (Mgmt), default<br>LAN 1 | 7<br>7 |
| RCST1 | SVN-0: 0x000002/8<br>SVN-1: 0x1000A1/8, tagged | VRF-0/VLAN-1<br>VRF-1/VLAN-1 | Internal (Mgmt)<br>LAN 1,default | 7<br>5 |
| RCST2 | SVN-0: 0x000003/8<br>SVN-1: 0x1000B1/8, tagged | VRF-0/VLAN-1<br>VRF-1/VLAN-1 | Internal (Mgmt)<br>LAN 1, VLAN-1 | 7<br>4 |

#### 8.2.2.2 Configuration example 2

Figure 8.4 shows configuration example 2, where two RCSTs belong to the same SVN (SVN-1). RCST1 supports three VLANs (VLAN-1, VLAN-2, VLAN-4) while RCST2 supports one VLAN (VLAN-1). An example configuration of the VLAN mapping is shown for this topology in Tables 8.4 and 8.5.

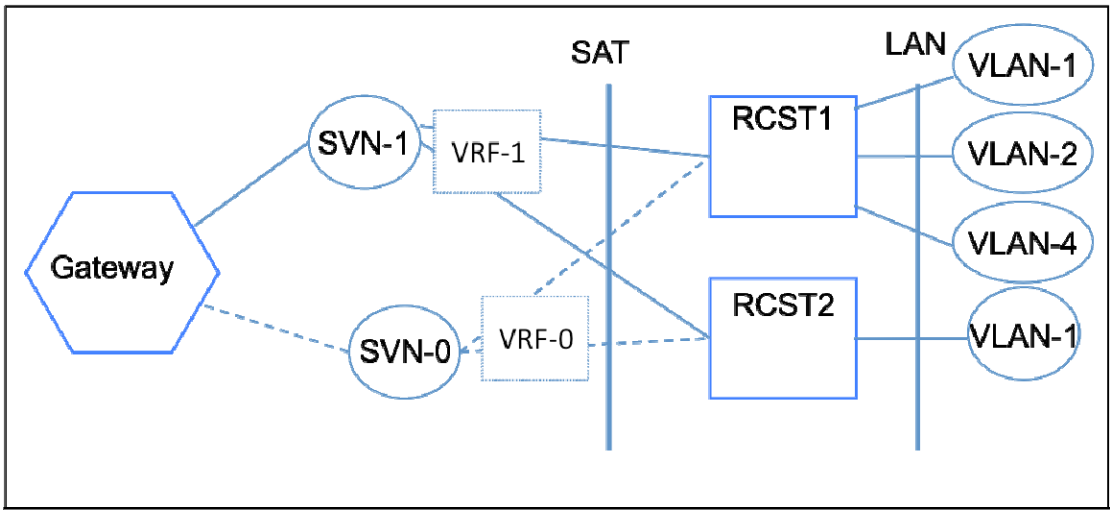


**Figure 8.4: Example 2: Two STs in one SVN; RCST1 supports three VLANs and RCST2 supports one VLAN**

In this example, Table 8.4 shows that the SVN_MASK of SVN-1 has a length of 12 bits, which allow the support of up to 4 095 SVNs.

If VLAN support is realized using method a), Table 8.4 shows an example configuration table. In this case, all tagged frames received at RCST1 from VLAN-2 and VLAN-4 will be encapsulated with a MAC24 of 0x1000A2 and 0x1000A3, respectively, in SVN-1 (prior removal of their tags). For untagged frames from VLAN-1, the MAC24 label of 0x1000A1 will be used.

NOTE: The RCST should discard any traffic that uses a VLAN_ID (or untagged default format) that is not explicitly listed in the table. In this case, the table does not configure RCST1 to support tagged frames with a VLAN_ID of 1, neither does it permit untagged frames to be forwarded by RCST2.

For RCST2, untagged frames will have a MAC24 of 0x1000B1 while tagged frames will be dropped.

**Table 8.4: Example of VLAN mapping to support SVN/VLAN using method a) for Figure 8.4 (default = with no tag on the Ethernet LAN interface; tagged = with a 802.1pQ tag)**

| | MAC24 | VLAN_ID | Interface |
|---|---|---|---|
| Gateway | SVN-0: 0x00000F/12<br>SVN-1: 0x100081/12<br>SVN-2: 0x100082/12<br>SVN-4: 0x100083/12 | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br>VRF-1/VLAN-2<br>VRF-1/VLAN-4 | LAN 0 (Mgmt)<br>LAN 1, VLAN-1<br>LAN 1, VLAN-2<br>LAN 1, VLAN-4 |
| RCST1 | SVN-0: 0x000001/12<br>SVN-1: 0x1000A1/12<br>SVN-1: 0x1000A2/12<br>SVN-1: 0x1000A3/12 | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br>VRF-1/VLAN-2<br>VRF-1/VLAN-4 | Internal (Mgmt)<br>LAN 1, default<br>LAN 1, VLAN-2<br>LAN 1, VLAN-4 |
| RCST2 | SVN-0: 0x000002/12<br>SVN-1: 0x1000B1/12 | VRF-0/VLAN-1<br>VRF-1/VLAN-1 | Internal (Mgmt)<br>LAN 1, default |

An example VLAN mapping to support of SVN/VLAN through method b) is shown in Table 8.5. This indicates that all tagged frames received at RCST1 from VLAN-2 and VLAN-4 will be encapsulated with a MAC24 of 0x1000A1 in SVN-1. If their PCP field is lower or equal to 7 and 5, respectively, there will be no changes to this value; otherwise the ST will set the PCP value to 4.Untagged frames will be tagged with a default VLAN_ID of VLAN-1 and, also, the MAC24 of 0x1000A1. This use makes an untagged frame equivalent to one with a VLAN_ID of VLAN-1 (although the current configuration does not expose this internal VLAN_ID on any LAN Interface. If their PCP field is lower or equal to 4, there will be no changes to this value; otherwise the ST will set the PCP value to 4. For RCST2, untagged frames will be also tagged with the VLAN_ID of VLAN-1 but a MAC24 of 0x1000B1 and, a maximum PCP of 4.

The Gateway is assumed to support 2 interfaces (0 and 1), used respectively for management and traffic, necessary because these use different VRF groups.

**Table 8.5: Example of VLAN mapping for Figure 8.4 for VLAN support (default = with no tag on LAN interface; tagged = with a 802.1pQ tag)**

| | MAC24 | VLAN_ID | Interface | PCP |
|---|---|---|---|---|
| Gateway | SVN-0: 0x00000F/12 | VRF-0/VLAN-1 | LAN 0 (Mgmt), default LAN 1 | 7 |
| | SVN-1: 0x100081/12, tagged | VRF-1/VLAN-1 | | 7 |
| | SVN-2: 0x100082/12, tagged | | LAN 1 | |
| | SVN-4: 0x100083/12, tagged | VRF-1/VLAN-2 | | 7 |
| | | | LAN 1 | |
| | | VRF-1/VLAN-4 | | 7 |
| RCST1 | SVN-0: 0x000001/12 | VRF-0/VLAN-1 | Internal (Mgmt) | 7 |
| | SVN-1: 0x1000A1/12, tagged | VRF-1/VLAN-1 | | 4 |
| | SVN-1: 0x1000A1/12, tagged | | LAN 1,default | |
| | SVN-1: 0x1000A1/12, tagged | VRF-1/VLAN-2 | | 7 |
| | | | LAN 1, VLAN-2 | |
| | | VRF-1/VLAN-4 | | 5 |
| | | | LAN 1, VLAN-4 | |
| RCST2 | SVN-0: 0x000002/12 | VRF-0/VLAN-1 | Internal (Mgmt) | |
| | SVN-1: 0x1000B1/12, tagged | VRF-1/VLAN-1 | LAN 1,default | 4 |

### 8.2.2.3 Configuration example 3

Figure 8.5 shows the case of two STs that each belong to a different SVN, SVN-1 and SVN-2 and each supporting one VLAN (VLAN-1 and VLAN-2, respectively).

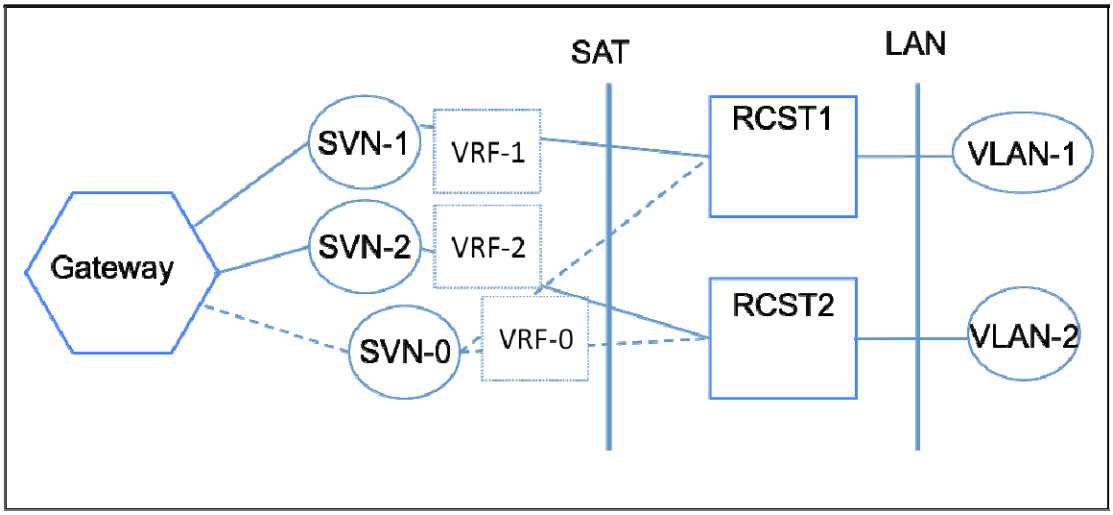


**Figure 8.5: Example 3: Two STs, each in one SVN and supporting one VLAN, respectively**

Table 8.6 shows an example VLAN Mapping for this topology using method a). In this case, the SVN_MASK has a length of 8 bits, which allows the support of up to 256 SVNs. Each ST appends a MAC24 corresponding to its VLAN_ID: 0x1000A1 for VLAN-1 (RCST1) and 0x1000B2 for VLAN-2 (RCST-2).

Any tagged frame arriving at an RCST will be dropped, because frames with any unassigned VLAN_ID cannot be forwarded (including the default in this case).

The Gateway uses a dedicated interface for management (0) and two traffic interfaces (1 and 2), since the traffic has been segregated into VRF groups. The Gateway interface 1 and 2 in this case enable the use of VLANs. (Since there is only one VLAN on each interface in this example this VLAN could have been mapped to a default interface with no VLAN tag).

**Table 8.6: Example of VLAN mapping for Figure 8.5 for VLAN support method a) (default = without tag on Ethernet LAN interface)**

| | MAC24 | VLAN_ID | Interface |
|---|---|---|---|
| Gateway | SVN-0: 0x00000F/8<br>SVN-1: 0x100081/8<br>SVN-2: 0x100082/8 | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br>VRF-2/VLAN-2 | 0 (Mgmt), default<br>1 VLAN-1<br>2 VLAN-2 |
| RCST1 | SVN-0: 0x000001/8<br>SVN-1: 0x1000A1/8 | VRF-0/VLAN-1<br>VRF-1/VLAN-1 | Internal (Mgmt)<br>LAN 1,default |
| RCST2 | SVN-0: 0x000002/8<br>SVN-2: 0x1000B2/8 | VRF-0/VLAN-1<br>VRF-2/VLAN-2 | Internal (Mgmt)<br>LAN 1, default |

Table 8.7 shows an example VLAN Mapping for this topology using method b). Each ST appends a MAC24 corresponding to its VLAN_ID: 0x1000A1 for VLAN-1 in SVN-1 (RCST1) and, 0x1000B1 for VLAN-2 in SVN-2 (RCST-2). The maximum PCP for the IEEE 802.1pQ [i.72] tag of packets arriving to RCST1 and RCST2 will set to 6 and 5, respectively.

RCST2 has been configured to transport any VLAN-ID that arrives at the LAN interface, carried within VRF-2. This utilizes the ability of method (b) to transport a VLAN_ID across the satellite link.

**Table 8.7: Example of VLAN mapping for Figure 8.5 for VLAN support method b) (default = with no tag on Ethernet LAN interface; tagged = with a 802.1pQ tag)**

| | MAC24 | VLAN_ID | Interface | PCP |
|---|---|---|---|---|
| Gateway | SVN-0: 0x00000F/8<br>SVN-1: 0x100081/8, tagged<br>SVN-2: 0x100082/8,<br>tagged | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br><br>VRF-2/any | 0 (Mgmt), default<br>1, tagged<br><br>2, tagged | 7<br>7<br><br>7 |
| RCST1 | SVN-0: 0x000001/8<br>SVN-1: 0x1000A1/8, tagged | VRF-0/VLAN-1<br>VRF-1/VLAN-1 | Internal (Mgmt)<br>LAN 1,tagged | 7<br>6 |
| RCST2 | SVN-0: 0x000002/8<br>SVN-2 0x1000B1/8, tagged | VRF-0/VLAN-2<br>VRF-2/any | Internal (Mgmt)<br>LAN 1,tagged | <br>5 |

### 8.2.2.4 Configuration example 4

Figure 8.6 shows a configuration example 4, where RCST1 supports two SVNs (SVN-1 and SVN-2) and two VLANs (VLAN-1, and VLAN-2). An example VLAN Mapping for this topology using method (a) and (b) to support VLANs is shown in Tables 8.8 and 8.9, respectively.

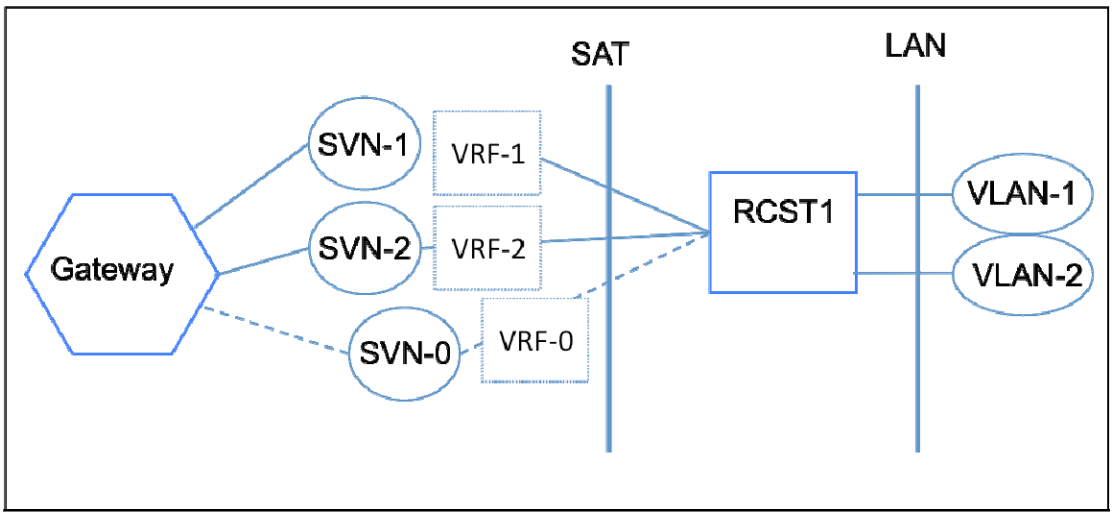


**Figure 8.6: Example 4: One ST supporting two SVNs and two VLANs**

Table 8.8 shows an example VLAN Mapping for this topology using method a). In this case, the SVN_MASK has a length of 10 bits, which allows the support of up to 1 023 SVNs. The ST uses a MAC24 label of 0x1000A1 to untagged packets (VLAN-1) while tagged packets are sent with a MAC24 of 0x1000A2 (VLAN-2).

The Gateway uses a dedicated interface for management (0) and a tagged interface for the VLANs (1). In this case, the RCST VLANs are mapped to new values at the egress interface 1 of the Gateway (e.g. SVN-1 to VLAN-4 and SVN-2 to VLAN-5), to allow the operator to differentiate this traffic from other VLANs configured within the network. This flexibility allowing remapping is common when VLANs are used.

**Table 8.8: Example of VLAN mapping to support SVN/VLAN using method a) for Figure 8.6 (default = with no tag on Ethernet LAN interface; tagged = with a 802.1pQ tag)**

| | MAC24 | VLAN_ID | Interface |
|---|---|---|---|
| Gateway | SVN-0: 0x0000F/10<br>SVN-1: 0x100081/10<br>SVN-2: 0x100082/10 | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br>VRF-2/VLAN-2 | 0 (Mgmt), default<br>1, VLAN-4<br>1, VLAN-5 |
| RCST1 | SVN-0: 0x000001/10<br>SVN-1: 0x1000A1/10<br>SVN-2: 0x1000A2/10 | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br>VRF-2/VLAN-2 | Internal (Mgmt)<br>LAN 1,default<br>LAN 1,tagged |

An example VLAN mapping using method b) is shown in Table 8.9. In this example, untagged frames, arriving to RCST1, are given a VLAN_ID of 1 and a maximum PCP of 3, corresponding to an MAC24 of 0x1000A1; while tagged frames will have a maximum PCP of 7 and a MAC24 of 0x1000A2.

**Table 8.9: Example of VLAN mapping to support SVN/VLAN using method b) for Figure 8.6 (default = without tag on Ethernet LAN interface; tagged = with a 802.1pQ tag)**

| | MAC24 | VLAN_ID | Interface | PCP |
|---|---|---|---|---|
| Gateway | SVN-0: 0x00000F/10<br>SVN-1: 0x100081/10, tagged<br>SVN-2: 0x100082/10,<br>tagged | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br><br>VRF-2/VLAN-2 | 0 (Mgmt), default<br>1, tagged<br><br>2, tagged | 7<br>7<br><br>7 |
| RCST1 | SVN-0: 0x000001/10<br>SVN-1: 0x1000A1/10, tagged<br>SVN-2: 0x1000A2/10, tagged | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br><br>VRF-2/VLAN-2 | Internal (Mgmt)<br>LAN 1,default<br><br>LAN1, tagged | <br>3<br><br>7 |

## 8.2.3 Multi-dwelling scenario

It is expected that multi-VLAN support at the LAN interface of the RCST will be attractive for multi-dwelling users. In this scenario, two or more subscribers share a terminal, but not necessarily the same QoS services. Each subscriber may use a different VLAN, mapped to a different SVN.

The presence of VLANs can provide isolation between different users (locations) connected to the multi-dwelling RCST LAN interface, (e.g. to support a VLAN switch including the ability to support additional untagged interface ports).

Multiple SVNs may be managed by the SVNO.

A use-case may support two sets of users via a single RCST, offering an independently managed SLA to each. At the RCST, both users are supported on a single LAN interface, through the use of dedicated VLANs. In this example, one uses an untagged VLAN and the second uses a tagged VLAN with a dedicated VLAN_ID value. The RCST may be connected to an external Ethernet switch that provides a dedicated (untagged) interface to the second user.

For this case, method a) is implemented: all frames arriving at the RCST are tagged. Table 8.10 shows an example VLAN Mapping for this topology. In this case, the SVN_MASK has a length of 10 bits, which allow the support of up to 1 023 SVNs. The RCST appends a SVN_MASK label of 0x1000A1(SVN-1) to packets with IP addresses corresponding to the VRF-1 group while packets are sent with a MAC24 of 0x1000A2 (SVN-2) if their IP addresses are from the VRF-2 group.

**Table 8.10: Example of VLAN mapping for SVN/VLAN support in multi-dwelling scenarios (default = with no tag on LAN interface)**

| | MAC24 | VLAN_ID | Interface |
|---|---|---|---|
| Gateway | SVN-0: 0x00000F/10<br>SVN-1: 0x100081/10<br>SVN-2:: 0x100082/10 | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br>VRF-2/VLAN-2 | 0 (Mgmt), default<br>1, default |
| RCST1 | SVN-0: 0x000001/10<br>SVN-1: 0x1000A1/10<br>SVN-2: 0x1000A2/10 | VRF-0/VLAN-1<br>VRF-1/VLAN-1<br>VRF-2/VLAN-2 | Internal (Mgmt)<br>LAN 1,default<br>LAN 2 |

### 8.2.4 SCADA scenario

This scenario will not typically support VLANs.

### 8.2.5 Backhauling scenario

For the backhauling scenario, VLAN support is not required and, usually, one SVN will be configured.

## 8.3 Recommendations for VLAN management

Some recommendations are provided in this clause for management of VLANs in interactive DVB-RCS2 networks. The proposal is based on the current MIB objects existing in [i.1]. However, a new table is needed in the MIB for mapping user VLANs and satellite SVNs.

### 8.3.1 Specifications of MIB objects

In the *interfaces* group of RCS2 MIB, there is an association of each interface with an MAC24 (parameter *ifPhysAddress*). One or several Ethernet interfaces may be used in the LAN of a RCST, each having its corresponding MAC24. Moreover, the same physical interface could correspond to several virtual (VLAN) interfaces.

In the *dvbRcs2NetworkConfig* group, the NetworkConfigTable associates each interface with its L3 network address. It supports the management interface and also the user interfaces. Note that every interface can be assigned an IPv4 or IPv6 address type. Parameter *NetworkConfigLANInetAddressIfIndex* is a link to the *interfaces* group table, therefore this table allows to configure all the virtual interfaces.

Table 8.11 is a new table that may be used to add VLAN support in the RCST. The objective of this table is to establish how to forward the VLAN frames received in the user interface. To achieve this, it is needed to map user VLAN and satellite SVNs (through the MAC24 address associated to each interface).

**Table 8.11: RCST MIB objects for VLAN mapping**

| Element | Type | Description |
|---|---|---|
| VLANmode | INTEGER | 0: Default mode. For packets in the ingress LAN interface, the RCST should remove the VLAN tag and encapsulate IP packets (when needed, depending on context) using a MAC24 with an SVN Mask derived from *VlanSvnMac* parameter in this table.<br>For packets in the egress interface, the RCST should tag the frame using the VLAN_ID associated to the MAC24 interface that has received the packet. Not tagging the frame for certain SVN interfaces is also possible, when *VlanId* = 0.<br><br>1: For packets in the ingress LAN interface, forward the IEEE 802.1pQ [i.72] tag through the satellite interface associated with the MAC24 address taken from the VLANMappingTable. Several or all of the VLAN_ID may be mapped to a single traffic MAC24 interface.<br>For packets received from the satellite interfaces, forward the received 802.1Q frames to the egress LAN interface. |
| VLANMappingTable | SEQUENCE OF VLANMappingTable ENTRY | Table that associates each VLAN_ID with an interface and set its properties. |
| VLANMappingTableEntry | SEQUENCE OF {<br>VlanInterfaceIndex,<br>VlanId,<br>VlanSvnMac,<br>VlanPcp,<br>VlanRowStatus} | |
| VlanInterfaceIndex | INTEGER | ST Interface number, that links to the *interfaces* group *ifNumber*. |
| VlanId | INTEGER | Corresponds to the 12-bit tag of a IEEE 802.1pQ [i.72] frame. |
| VlanMAC24 | OCTET STRING | The only possible values for this parameter are the values populated in *L3VirtualRoutingForwardingConfig* group, obtained during RCST logon.<br>For outgoing frames, this parameter is the MAC24 address of the satellite interface that will be used when the Ethernet frame comes with VLAN_ID equal to *VlanId* value of the same row. Untagged frames can also be mapped to a certain MAC24.<br>For packets from the Satellite (egress interface), the VLAN_ID to be tagged by the RCST will depend on the MAC24 of the received PPDU frame. The SVN Mask of the received frame is used to infer the VLAN_ID tag. |
| VlanPcp | | Maximum priority code point. A higher value of the PCP in the IEEE 802.1pQ [i.72] frame will be decremented to this value. Applicable for VLANmode = 1.<br>This value is used by the ST QoS model for the PHB association with the VLAN. |
| VlanRowStatus | Row Status | The row status, used according to row creation and removal conventions. A row entry cannot be modified when the status is marked as active(1). A row can be created either by createAndGo and automatically change to active state or createAndWait to add more parameters before becoming active. |

# 9 PEP session negotiation protocol

## 9.1 State definitions

The following states (illustrated in Figures 9.1 and 9.2) are relevant in the context of PEP Session Negotiation:

**Off/Standby:** This is the normal state immediately following power-on initialization, as well as a default state to which the RCST returns in some situations following loss of synchronization or upon being logged off. It is an implementation choice whether this state is absorbing; i.e. whether any external stimulus is required in order to initiate the processes that may cause a transition away from this state. The forward link should be kept operational in this state. When entering the Off/Standby state, the RCST should immediately cease transmission. It may keep dynamic identifiers if specifically allowed to do so as indicated for the assignment. The RCST should not transmit while in the Off/Standby state.

**Hold/Standby:** When entering the Hold/Standby state, the RCST should immediately cease transmission. It may keep dynamic identifiers if specifically allowed to do so as indicated for the assignment. An RCST in the Hold/Standby state should remain there following restart and power cycling events until the NCC releases the conditions(s) that keep the RCST in the Hold/Standby state. The forward link should be kept operational in this state. The RCST should not transmit while in the Hold/Standby state.

**Ready for Logon:** The RCST enters this state when the forward link has been successfully acquired and the configuration data required for issuing logon is up to date. It is an implementation choice whether this state is absorbing; i.e. whether any external stimulus is required to initiate the processes that may cause a transition away from this state. External triggers may include for example arrival of data on the terrestrial interface or reception of a "wake-up" message in the TIM-U. Transmission of logon bursts is allowed when the RCST is in this state.

**PEP Advertise Received:** The RCST enters this state upon reception of the pep_control_advertise message. This message can be broadcast any time, or received right after the logon process. Upon reception of the message, the RCST can either send a pep_control_offer or reply with an error message aborting the process.

**Offering PEP:** This state is entered once the pep_control_offer message has been sent by the RCST, This message is sent either upon reception of the pep_control_advertise message, or any time to force renegotiation of the PEP to be used for a given active SVN-MAC The terminal will wait for a response from the hub (pep_control_use), and if not received, after a number of retries and t/o expired, will abort the process.

**PEP Use Received:** The RCST enters this state upon reception of the pep_control_use message that instructs the RCST to use one of the offered PEPs for the SVN-MAC on which it is received. A PEP Control Use Message may be sent at any time for any active SVN-MAC. In case the RCST cannot activate the required PEP configuration, it should return an error code to report the problem. Otherwise it will automatically transit to the next state.

**Use PEP on MACx:** Successful processing of a PEP Control Use message causes the RCST to enter this state and to use the instructed PEP for the SVN-MAC on which it is received. PEP negotiation completion causes a transition away from this state.

**TDMA Sync:** This is the normal operational state for the RCST. This is an absorbing state; the RCST should remain there until external events or loss of TDMA synchronization dictate transition to another state. The TDMA synchronization status should be supervised by the Sync Monitoring Process. Transmission of control bursts is allowed when the RCST is in this state. Transmission of traffic burst and traffic/control bursts may be allowed or these may be dynamically blocked even if assigned.

**NCR Recovery:** The RCST enters this state when there is loss of TDMA synchronization or NCR loss when in TDMA Sync. This is a non-absorbing state; the RCST should autonomously transition to another state. The RCST should not transmit while in the NCR Recovery state.

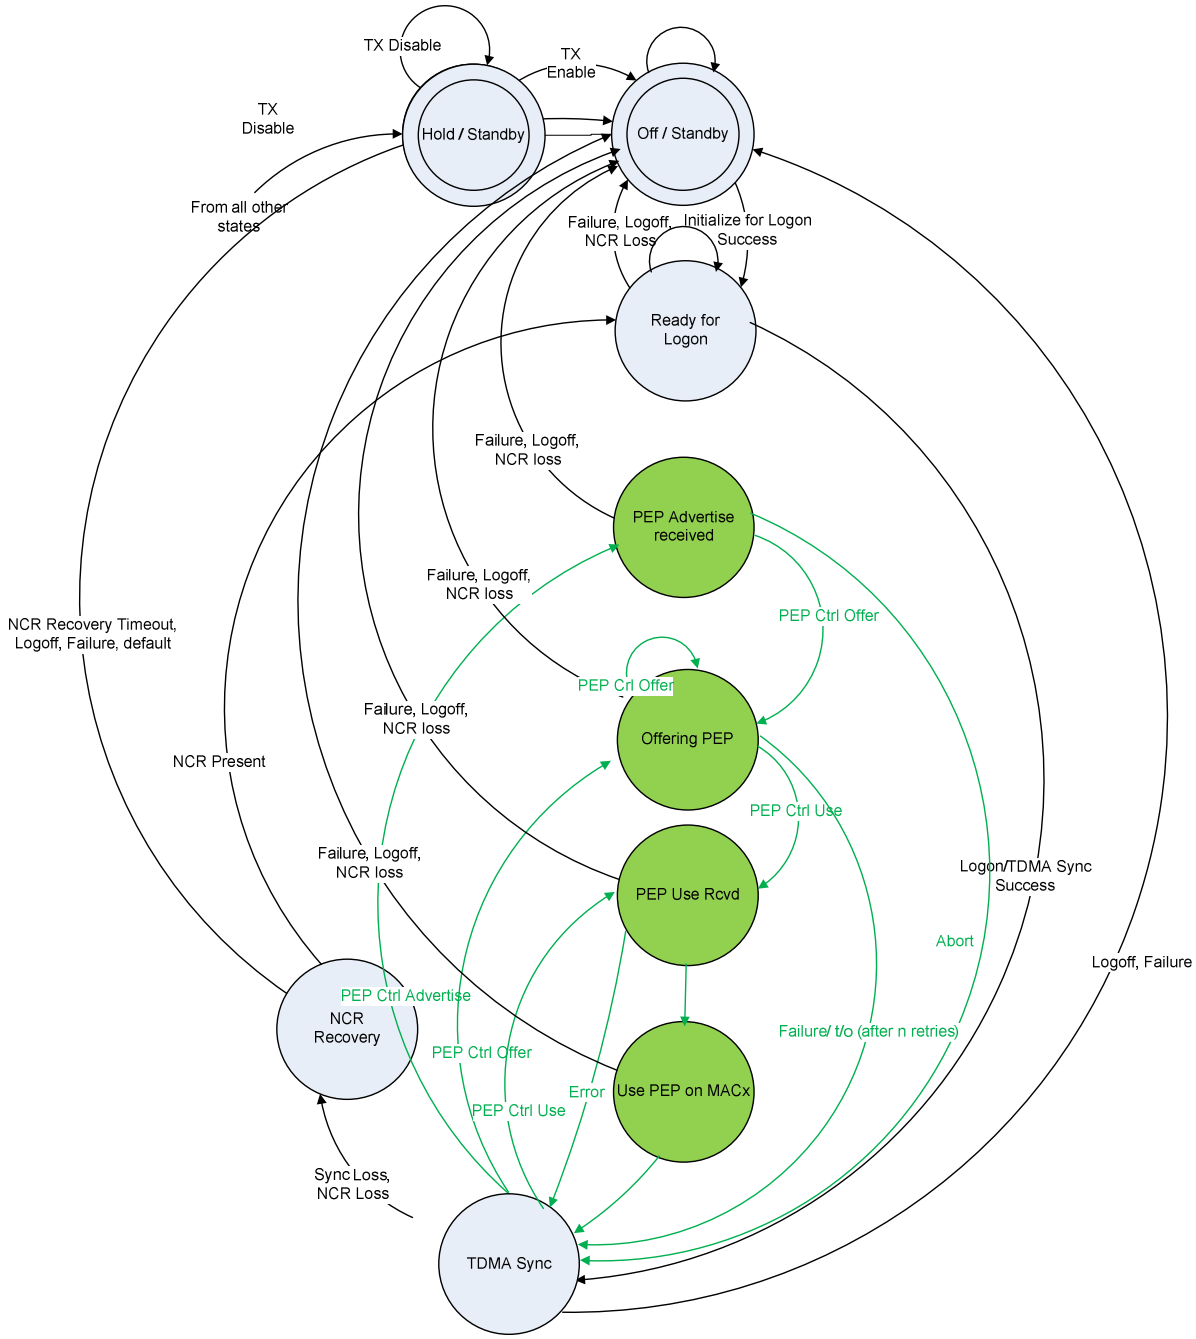


**Figure 9.1: RCST State Transition Diagram for PEP Session Negotiation Protocol**

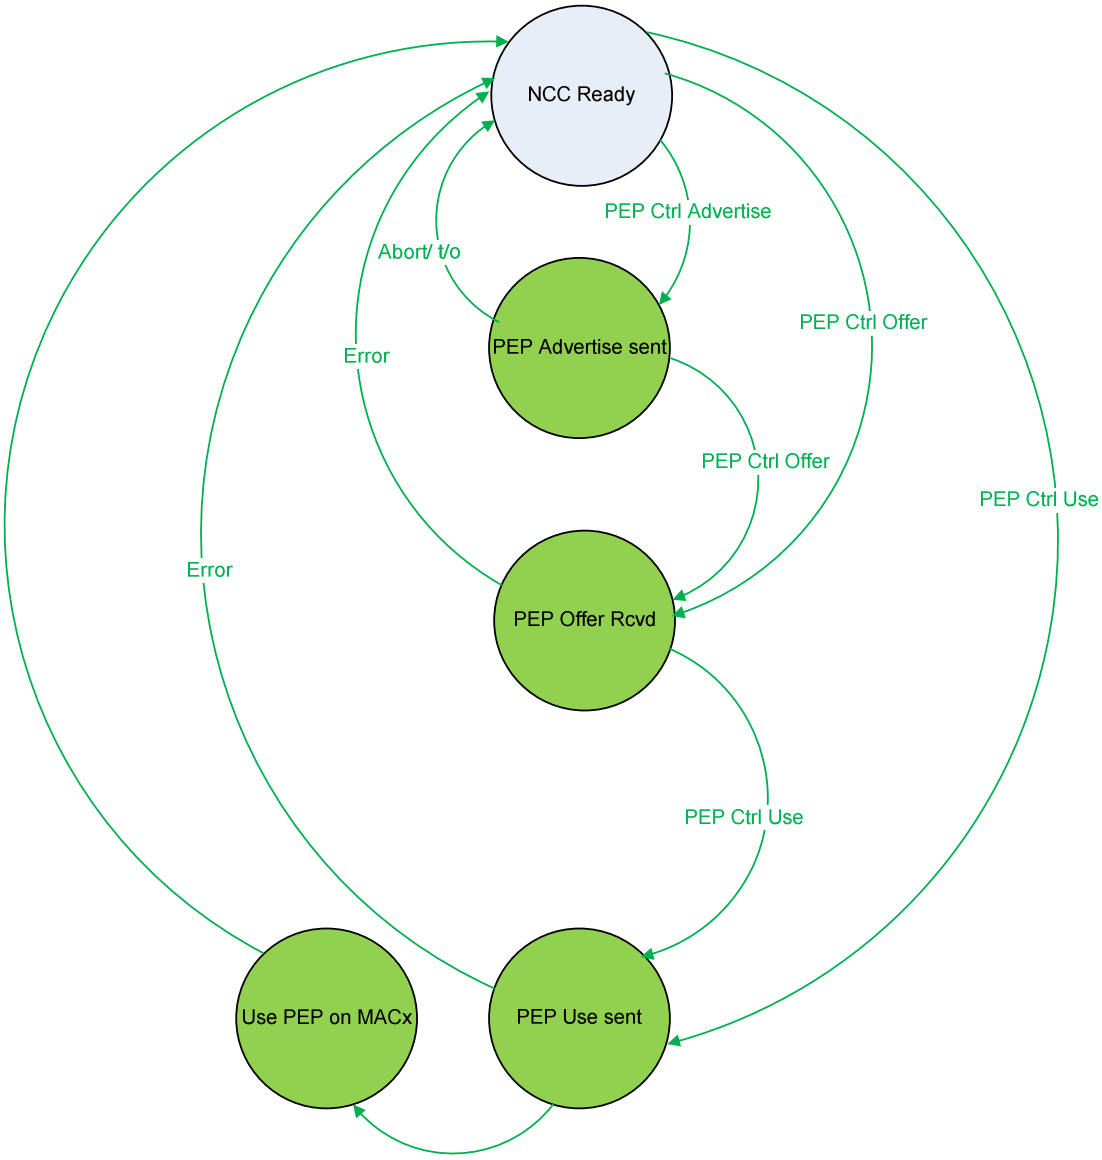


**Figure 9.2: Hub State Diagram for PEP Session Negotiation**

## 9.2 PEP negotiation protocol parameters and MIB group

The PEP negotiation protocol makes use of the HLS agent control protocol (clause 9.2 in [i.1]). This protocol is used over the IPv4 address provisioned for a satellite interface and bound to a MAC24 label for management signalling.

The PEP negotiation group in the RCST MIB from [i.1] compiles all the necessary information to perform PEP negotiation between the RCST and the NCC.

## 9.3 Example use cases

An example of the message exchange during a normal progression of PEP negotiation is illustrated in Figure 9.3. The sequence illustrates the normal flow of events and signals.

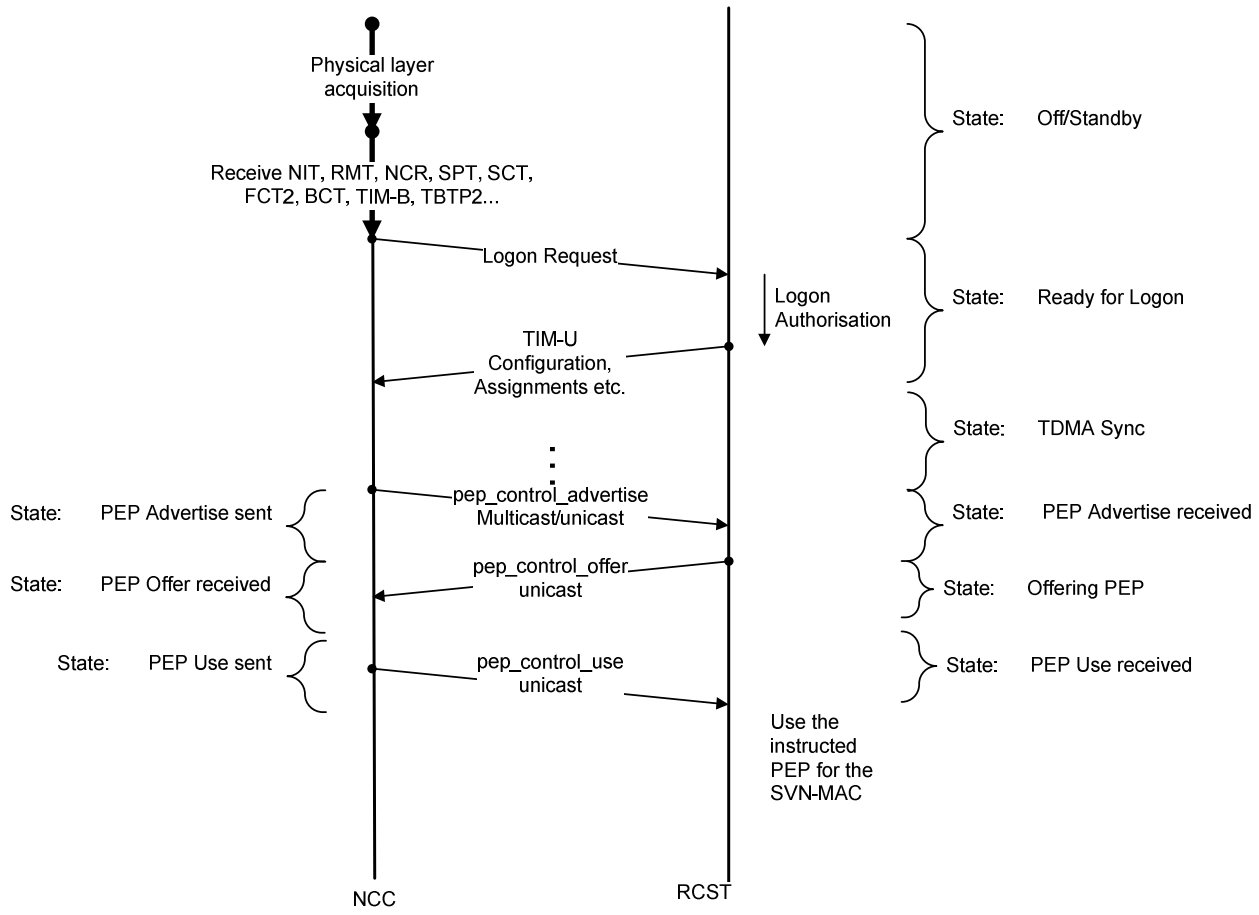


**Figure 9.3: Normal PEP negotiation**

The RCST supports the current set of messages for TCP-PEP negotiation [i.1]. Each offer contains N descriptors for the offered TCP-PEPs. Each response contains M descriptors for the supported TCP-PEPs, where M=<N. The NCC finally selects one TCP-PEP.

The transport of RCST Agent negotiation messages is explained below:

1) The IPv4 multicast group destination address and UDP port number are received via HLID descriptor in the TIM-U.

2) A PEP Advertise message is received on the forward link. This forward IP message is either unicast to the RCST IPv4 address or multicast to the multicast group address in step-1. The destination UDP port number for this forward IP message is as in step-1.

3) RCST sends a PEP Offer message with a destination IPv4 address that matches the IP source address of the PEP Advertise message and using the UDP destination port that was used in the PEP Advertise message. The IP packet is sent with the IP source of the RCST and using the same SVN on which the PEP Offer was received.

4) A PEP Use or PEP Error message is sent in response to a PEP Offer message. This has an IP source address that is identical to the IP destination address of the PEP Offer and a IPv4 destination address identical to the IP source address used for the PEP Offer. The UDP source port is identical to the UDP destination port of the PEP Offer message.

The above exchange is used to configure the PEP used for a specific SVN. An RCST that supports multiple SVNs should repeat steps 3 & 4 of this negotiation for each SVN that is active.

### 9.3.1 Consumer/SCADA/Backhauling scenarios

For these scenarios, an RCST will likely be part of only one traffic SVN. Figure 9.4 illustrates the message exchange that corresponds to this scenario.

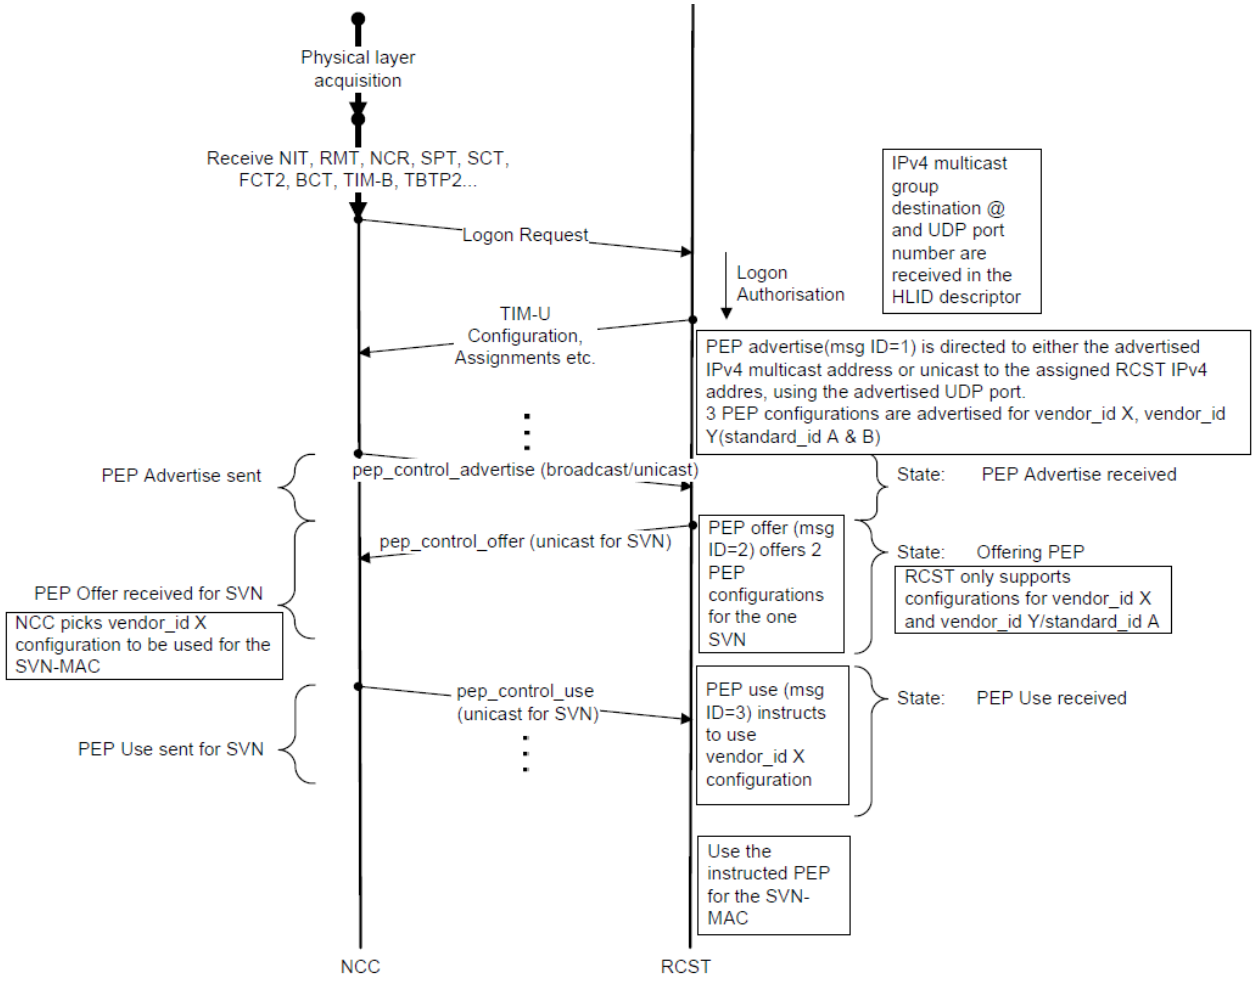


**Figure 9.4: Normal PEP negotiation for one SVN**

### 9.3.2 Corporate/Institutional/Multi-dwelling scenarios

The RCST may be part of one or more traffic SVNs It is expected that for multi-dwelling users multiple SVNs may be managed by the OVN. Steps 3 & 4 of the negotiation are repeated for each active SVN. In the next example the RCST issues a PEP Control Offer Message for two of its active MAC24s. The offer forces renegotiation of the PEP to be used for the MAC24s.

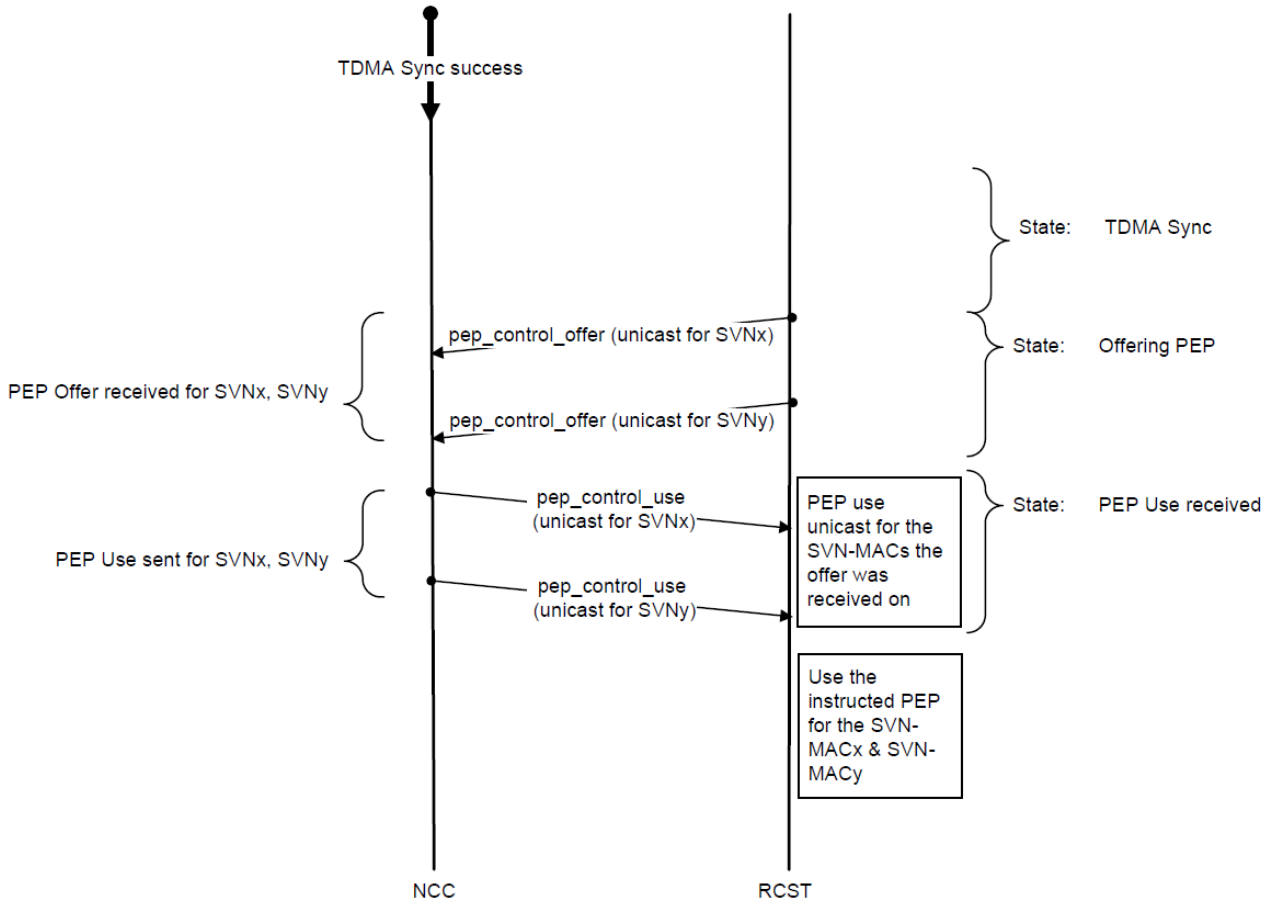


**Figure 9.5: Normal PEP negotiation for two SVNs**

# 10 SNMP configuration

The recommended management reference network for RCS follows the TMN model of telecom network management to help the operators to configure and manage the RCS network in an easy way. In this architecture, the NMC performs all management functions, namely system configuration, fault management, system performances management and accounting data retrieval (FCAPS functions). The NMC and NCC could either be directly connected through a LAN interface, or via IP connection over terrestrial backhaul networks. The basic functionality of the NMC includes the manager of the elements of the network (RCST, GW, NCC). These functions support a SNMPv2c/SNMPv3 protocol and MIB data base (in the communication between NMC and network elements - Internal interface). The NMC is the SNMP manager and the RCST, NCC or Gateway are the SNMP agents.

To comply with the recommended management architecture, the RCST will require a default or minimum SNMP configuration before a successful logon. This data should be provided by the installer or first configuration file.

This clause provides the default and operational SNMP configuration for the different management actors/roles in the network.

The RCST may use the following tables to provide the desired SNMP Access:

- snmpCommunityTable [i.32] for SNMP community configuration
- snmpTargetAddrTable [i.33]
- snmpTargetAddrExtTable [i.32]: The table of mask and maximum message size (mms) value associated with the snmpTargetAddrTable
- vacmAccessTable [i.34]: view access table configuration

Access to an SNMP server by an SNMP client is governed by a proprietary SNMP community table that identifies those communities that have access to MIB data.

When an SNMP server receives a request, the server extracts the client's IP address and the community name. The SNMP community table is searched for a matching community. If a match is found, its access list name is used to validate the IP address. If the access list name is null, the IP address is accepted. A nonmatching community or an invalid IP address results in an SNMP authentication error.

Each entry in the community table identifies:

- SNMP community name: public / private or a new name
- SNMP community security name: A human readable string representing the corresponding value of snmpCommunityName in a Security Model independent format
- snmpCommunityContextEngineID
- snmpCommunityContextName
- snmpCommunityTransportTag: This object specifies a set of transport endpoints from which a command responder application will accept management requests. If a management request containing this community is received on a transport endpoint other than the transport endpoints identified by this object, the request is deemed unauthentic. The transports identified by this object are specified in the snmpTargetAddrTable. Entries in that table whose snmpTargetAddrTagList contains this tag value are identified. If the value of this object has zero-length, transport endpoints are not checked when authenticating messages containing this community string

For a first default SNMP configuration, it is recommended to have only public / private communities, and to ensure a minimum level of protection only with the IP address of the primary NMC and mask 255.255.255.0. The default communities can be changed or additional ones can be added.

The View Based Access Control Model (VACM) from [i.34] defines the necessary elements of procedure for controlling access to management information.

To implement the View Based Access Control Model (VACM) an SNMP entity needs to retain information about access rights and policies. This information is part of the SNMP engine's Local Configuration Datastore (LCD). See [i.35] for the definition of LCD. In order to allow an SNMP entity's LCD to be remotely configured, portions of the LCD need to be accessible as managed objects. A MIB module, the View-based Access Control Model Configuration MIB, defines these managed object types.

Figure 10.1 shows how the decision for access control is made by the view based access control model:

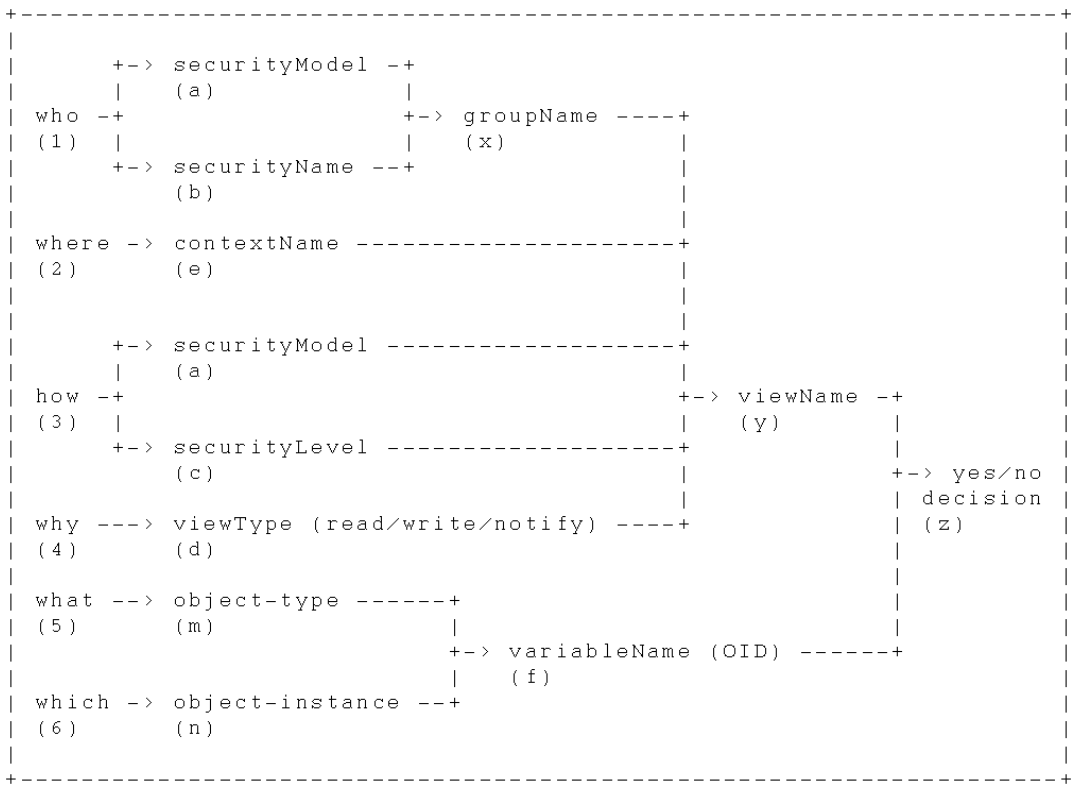


**Figure 10.1: Access control decision by VACM**

How the decision for isAccessAllowed is made:

1) Inputs to the isAccessAllowed service are:

   a) securityModel -- SNMPv3 was designed for the use of multiple co-existing security models. The msgSecurityModel field specifies the security model that was used to generate the message. Therefore, the receiving entity knows which security model should be used to perform security processing upon message reception

   b) securityName -- principal who wants to access (as specified in the community table)

   c) securityLevel -- Level of Security: Different access rights for members of a group can be defined for different levels of security, i.e. noAuthNoPriv, authNoPriv, and authPriv. The securityLevel identifies the level of security that will be assumed when checking for access rights (see the SNMP Architecture document [i.35] for a definition of securityLevel). The View-based Access Control Model requires that the security Level is passed as input to the Access Control module when called to check for access rights.

   d) viewType -- read, write, or notify view

   e) contextName -- context containing variableName

   f) variableName -- OID for the managed object

      -- this is made up of:

      - object-type (m)

      - object-instance (n)

2) The partial "who" (1), represented by the securityModel (a) and the securityName (b), are used as the indices (a,b) into the vacmSecurityToGroupTable to find a single entry that produces a group, represented by groupName (x).

3) The "where" (2), represented by the contextName (e), the "who", represented by the groupName (x) from the previous step, and the "how" (3), represented by securityModel (a) and securityLevel (c), are used as indices (e,x,a,c) into the vacmAccessTable to find a single entry that contains three MIB views.

4) The "why" (4), represented by the viewType (d), is used to select the proper MIB view, represented by a viewName (y), from the vacmAccessEntry selected in the previous step. This viewName (y) is an index into the vacmViewTreeFamilyTable and selects the set of entries that define the variableNames which are included in or excluded from the MIB view identified by the viewName (y).

5) The "what" (5) type of management data and "which" (6) particular instance, represented by the variableName (f), is then checked to be in the MIB view or not, e.g. the yes/no decision (z).

As an example, the VACM configuration for SNO and SVNO basic access roles would be:

The initial parameters that should be configured during installation for the View-based Access Control Model are:

**A security configuration**: The choice of security configuration determines if initial configuration is implemented and if so how. One of three possible choices is selected:

- initial-minimum-security-configuration
- initial-semi-security-configuration
- initial-no-access-configuration

In the case of a initial-no-access-configuration, there is no initial configuration, and so the following steps are irrelevant.

6) Community table: Three entries in the snmpCommunityTable, "initial", "sno", & "svno"

7) A default context: One entry in the vacmContextTable with a contextName of "" (the empty string), representing the default context. Note that this table gets created automatically if a default context exists.

vacmContextName ""

8) An initial group: One entry in the vacmSecurityToGroupTable to allow access to group "initial".

vacmSecurityModel 3 (USM)

vacmSecurityName "initial"

vacmGroupName "initial"

vacmSecurityToGroupStorageType anyValidStorageType

vacmSecurityToGroupStatus active

A SNO, and SVNO groups:

vacmSecurityModel 3 (USM)

vacmSecurityName "sno"

vacmGroupName "sno"

vacmSecurityToGroupStorageType anyValidStorageType

vacmSecurityToGroupStatus active

vacmSecurityModel 3 (USM)

vacmSecurityName "svno"

vacmGroupName "svno"

vacmSecurityToGroupStorageType anyValidStorageType

vacmSecurityToGroupStatus active

9) Initial access rights: Three entries in the vacmAccessTable as follows:

- read-notify access for securityModel USM, securityLevel "noAuthNoPriv" on behalf of securityNames that belong to the group "initial" to the <restricted> MIB view in the default context with contextName "".
- read-write-notify access for securityModel USM, securityLevel "authNoPriv" on behalf of securityNames that belong to the group "svno" to the <SNO> MIB view in the default context with contextName "".
- read-write-notify access for securityModel USM, securityLevel "authNoPriv" on behalf of securityNames that belong to the group "sno" to the <SNO> MIB view in the default context with contextName "".
- if privacy is supported, read-write-notify access for securityModel USM, securityLevel "authPriv" on behalf of securityNames that belong to the group "sno" to the <SNO> MIB view in the default context with contextName "".
- That translates into the following entries in the vacmAccessTable.
- One entry to be used for unauthenticated access (noAuthNoPriv):

vacmGroupName "initial"

vacmAccessContextPrefix ""

vacmAccessSecurityModel 3 (USM)

vacmAccessSecurityLevel noAuthNoPriv

vacmAccessContextMatch exact

vacmAccessReadViewName "restricted"

vacmAccessWriteViewName ""

vacmAccessNotifyViewName "restricted"

vacmAccessStorageType anyValidStorageType

vacmAccessStatus active

- Two entries to be used for authenticated access (authNoPriv) with optional privacy (authPriv):

vacmGroupName "svno"

vacmAccessContextPrefix ""

vacmAccessSecurityModel 3 (USM)

vacmAccessSecurityLevel authNoPriv

vacmAccessContextMatch exact

vacmAccessReadViewName "SVNO"

vacmAccessWriteViewName "SVNO"

vacmAccessNotifyViewName "SVNO"

vacmAccessStorageType anyValidStorageType

vacmAccessStatus active

vacmGroupName "sno"

vacmAccessContextPrefix ""

vacmAccessSecurityModel 3 (USM)

vacmAccessSecurityLevel authNoPriv

vacmAccessContextMatch exact

vacmAccessReadViewName "SNO"

vacmAccessWriteViewName "SNO"

vacmAccessNotifyViewName "SNO"

vacmAccessStorageType anyValidStorageType

vacmAccessStatus active

10) Two MIB views, of which the second one depends on the security configuration.

- Two views, the <SNO> view, and the <SVNO> for authenticated access:
  - the <SNO> MIB view is the following subtree: "internet" (subtree 1.3.6.1)
  - the <SVNO> MIB view is the following subtree: "internet" (subtree 1.3.6.1)
- A second view, the <restricted> view, for unauthenticated access. This view is configured according to the selected security configuration:
  - For the initial-no-access-configuration there is no default initial configuration, so no MIB views are prescribed.
  - For the initial-semi-secure-configuration:

  the <restricted> MIB view is the union of these subtrees:

  (a) "system" (subtree 1.3.6.1.2.1.1) [i.36]

  (b) "snmp" (subtree 1.3.6.1.2.1.11) [i.36]

  (c) "snmpEngine" (subtree 1.3.6.1.6.3.10.2.1) [i.35]

  (d) "snmpMPDStats" (subtree 1.3.6.1.6.3.11.2.1) [i.37]

  (e) "usmStats" (subtree 1.3.6.1.6.3.15.1.1) [i.38]
- For the initial-minimum-secure-configuration:

  the <restricted> MIB view is the following subtree.

  "internet" (subtree 1.3.6.1)

This translates into the "SNO" and "SVNO" entries in the vacmViewTreeFamilyTable.

vacmViewTreeFamilyViewName "SNO

vacmViewTreeFamilySubtree 1.3.6.1

vacmViewTreeFamilyMask ""

vacmViewTreeFamilyType 1 (included)

vacmViewTreeFamilyStorageType anyValidStorageType

vacmViewTreeFamilyStatus active

vacmViewTreeFamilyViewName "SVNO
vacmViewTreeFamilySubtree 1.3.6.1
vacmViewTreeFamilyMask ""
vacmViewTreeFamilyType 1 (included)
vacmViewTreeFamilyStorageType anyValidStorageType
vacmViewTreeFamilyStatus active

| | minimum-secure | semi-secure |
|---|---|---|
| vacmViewTreeFamilyViewName | "restricted" | "restricted" |
| vacmViewTreeFamilySubtree | 1.3.6.1 | 1.3.6.1.2.1.1 |
| vacmViewTreeFamilyMask | "" | "" |
| vacmViewTreeFamilyType | 1 (included) | 1 (included) |
| vacmViewTreeFamilyStorageType | anyValidStorageType | anyValidStorageType |
| vacmViewTreeFamilyStatus | active | active |

vacmViewTreeFamilyViewName "restricted"
vacmViewTreeFamilySubtree 1.3.6.1.2.1.11
vacmViewTreeFamilyMask ""
vacmViewTreeFamilyType 1 (included)
vacmViewTreeFamilyStorageType anyValidStorageType
vacmViewTreeFamilyStatus active

vacmViewTreeFamilyViewName "restricted"
vacmViewTreeFamilySubtree 1.3.6.1.6.3.10.2.1
vacmViewTreeFamilyMask ""
vacmViewTreeFamilyType 1 (included)
vacmViewTreeFamilyStorageType anyValidStorageType
vacmViewTreeFamilyStatus active

# 11 Terminal start-up phases

The objective of this clause is to show, step by step, the necessary functions, messages, and parameters required for the successful operation of a terminal in an RCS2 network; starting from the installation of the terminal and from there reach the operational status, understanding this status as the stage when the terminal is able to receive and transmit traffic.

This analysis aims at putting together concepts coming from the LL [i.3] and HL [i.1] specification and going into the fine details of the messages used and values of the parameters exchanged.

The following phases will be analyzed, from a first RCST power up to a successful network entry:

- RCST installation
- RCST forward link alignment
- RCST return link alignment
- RCST logon and first commissioning

Three different actors may perform M&C operations on the RCST:

- The SNO: responsible of RCST forward and return alignment and first logon into the network. The SNO is responsible for organizing the RCSTs in different Group_Ids and Logon_Ids and registering the non-volatile RCST HW addresses. Each RCST is given an SVN-MAC that can be used for management and control traffic from the SNO.
- The SVNO: responsible for the RCST traffic functions, IP routing, QoS, etc. The SVNO would be considered with a role of ISP with management functions and access to NMC client. The system profile parameters are set by the SVNO. The SNO assigns a set of SVN-MACs per SVNO. The SVNO is responsible of the distribution of a given sub-set of SVN-MACs between its SVNs. One or more SVN-MACs may be assigned to each RCST for traffic interfaces.
- The installer: responsible of the first set up of the installation parameters required for the RCST start up. The minimum set of parameters provided by the installer should be:
  - Operational forward link acquisition parameters
  - SNMP parameters for remote SNMP communication between the RCST and the SNO, and local SNMP from the installer
  - Fwd and Rtn alignment parameters (in case alignment is required)
  - Some of the System parameters (see table in clause 11.1.2)

This clause will conclude on what are the parameters that should be remotely accessed by the remote management entities and how.

## 11.1 RCST installation

After an RCST power on, and before connecting the RCST to the Operator Virtual Network (OVN), the RCST should count with an initial set of configuration parameters for the start-up. This set of parameters would allow the RCST to acquire the forward link, unless a pointing alignment process is required.

Once the forward link is acquired, the combination of ONID (Original Network ID) and INID will determine the SNO domain where the RCST belongs to. In terms of RCST operation, the SNO domain is transparent to the RCST.

### 11.1.1 Forward link acquisition parameters

The minimum set of parameters needed at initial installation for forward link acquisition is:

- ODU parameters within the System group (as already included in [i.1])
- Flink configuration group set of parameters (as already included in [i.1])

The ODU parameters use the same format that is used in SatLabs MIB [i.39]. By default they are considered RW only for the Installer. Anyhow they are recommended to be RW parameters for the SNO also, to allow remote configuration in case there is any problem.

**Table 11.1: ODU parameters**

| Functional Group | dvbRcs2SystemConfig | | | | | |
|---|---|---|---|---|---|---|
| **Element** | **Type** | **Unit** | **Range** | **Description** | **Access Rights** | **M&C Actor** |
| dvbRcs2SystemOduAntennaSize | INTEGER32 | cm | - | Diameter of the antenna. For supervision. | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision |
| dvbRcs2SystemOduSspa | INTEGER32 | 0,1 W | - | Power level of the Solid State Power Amplifier. For supervision. | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision |
| dvbRcs2SystemOduGain | INTEGER32 | 0,1 dBi | - | Antenna peak gain of the ODU. For supervision. | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision |
| dvbRcs2SystemOduTxType | SnmpAdmin String | | - | Type of transmitter installed in the ODU. For supervision. | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision |
| dvbRcs2SystemOduRxType | SnmpAdmin String | | - | Type of LNB installed in the ODU, with information such as vendor type, output type. For supervision. | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision, only in RCST MIB |
| dvbRcs2SystemOduRxBand | INTEGER | | High-band (0), Low Band (1) | LNB high band / Low band selector. High band corresponds to the emission of an 18-26 kHz tone with 0,4-0,8 Vpp in the Rx IFL cable. **Required for forward link acquisition.** | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision, only in RCST MIB |
| dvbRcs2SystemOduRxLO | INTEGER32 | | - | ODU reception local oscillator frequency. **Required for forward link acquisition.** | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision, only in RCST MIB |
| dvbRcs2SystemOduTxLO | INTEGER32 | In 100 Hz | - | ODU transmission Frequency of Block Up-Converter Local Oscillator. **Required for forward link acquisition.** | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision, only in RCST MIB |

In [i.1], the Flink configuration group lists the forward link attachment points (e.g. different for installation and operation), in a similar way that was done in SatLabs MIB. This table describes the forward link parameters used for the start up stream of the NCC as the follows:

- fwdStartPopId: population ID associated with the start up forward link.
- fwdStartFrequency: frequency of the start transponder carrying a NIT to which any RCST should trigger to acquire forward link.
- fwdStartPolar: polarization of the start transponder carrying the NIT.

- fwdStartFormat: transmission format standard applied to the start up stream (only dvbs2ccm or dvbs2acm would be allowed).
- fwdStartRolloff: roll-off applied on the start transponder (0.10, 0.20, 0.25, 0.35). Noted that in RCS2, the additional value of 0.10 has been added.
- fwdStartSymbolRate: symbol rate on the start transponder carrying a NIT to which any RCST should trigger to acquire the forward link information.
- fwdStartInnerFec: specifies the inner FEC on the start transponder.

Only the fwdStartPopId (operational population ID), fwdStartFrequency (fwd link frequency), fwdStartPolar and fwdStartFormat are really required to acquire the forward link. The rest of parameters can be used to check the fwd link being acquired. If no match is produced the RCST could give a warning.

The following parameters are also part of the Satellite Forward Link Descriptor [i.3]:

- Polarization
- Format
- RollOff
- SymbolRate
- InnerFEC

The Fwd Link Descriptor includes as well the satellite ID, beam ID, NCC ID, & local_multiplex ID that can be correlated with the Population ID through the RMT.

The set of parameters already included in RCS2 MIB, under Flink configuration group, are based on SatLabs MIB.

The Flink configuration parameters should be set by the installer in accordance to the RCST provisioning information kept in the SNO. They are recommended to be RW to allow remote reconfiguration in case there is any problem or provisioning change in the SNO for that particular RCST (e.g. change to a different frequency or coverage area).

**Table 11.2: FLink config parameters**

| Functional Group | dvbRcs2FwdConfiguration | | | | | |
|---|---|---|---|---|---|---|
| **Element** | **Type** | **Unit** | **Range** | **Description** | **Access Rights** | **M&C Actor** |
| dvbRcs2FwdStartEntry | SEQUENCE { dvbRcs2FwdStartIndex, dvbRcs2FwdStartPopID, dvbRcs2FwdStartFrequency, dvbRcs2FwdStartPolar, dvbRcs2FwdStartFormat, dvbRcs2FwdStartRolloff, dvbRcs2FwdStartSymbolRate, dvbRcs2FwdStartInnerFec, dvbRcs2FwdStartRowStatus | | | Fwd link configuration parameters | Installer (RW) SNO (RW) SVNO (RO) | SNO provisioning parameters as part of RCST profile |

The fwdStatus lists all the forward link status parameters, as RO parameters, for supervision. This group provides details on the forward link that the RCST has attached to. Right now this set of parameters is provided in the State group of RCS2 MIB.

**Table 11.3: FLink status parameters**

| Functional Group | dvbRcs2State | | | | | |
|---|---|---|---|---|---|---|
| **Element** | **Type** | **Unit** | **Range** | **Description** | **Access Rights** | **M&C Actor** |
| dvbRcs2FwdLinkStatus | INTEGER | | (0) notAcquired, (1) acquired | Provides the status of the RCST forward link. | RO | |
| dvbRcs2FwdStatusEntry | SEQUENCE {dvbRcs2FwdStatusIndex, dvbRcs2FwdStatusIfReference, dvbRcs2FwdStatusONetId, dvbRcs2FwdStatusNetId, dvbRcs2FwdStatusNetName, dvbRcs2FwdStatusFormat, dvbRcs2FwdStatusFrequency, dvbRcs2FwdStatusPolar, dvbRcs2FwdStatusInnerFec, dvbRcs2FwdStatusSymbolRate, dvbRcs2FwdStatusRolloff, dvbRcs2FwdStatusModulation, dvbRcs2FwdStatusFecFrame, dvbRcs2FwdStatusPilot, dvbRcs2FwdStatusBer, dvbRcs2FwdStatusCnr, dvbRcs2FwdStatusRxPower} | | | An entry in the forward link status table. Each entry is associated with a physical interface. | RO | |

## 11.1.2 RCST system parameters

As already introduced in the Management clause of [i.1] the RCST system profile is given by a set of parameters. These parameters are grouped as follows:

- System profile map (Consumer, SOHO, Multi-dwelling, corporate, SCADA, Backhaul, Institutional) that identifies the terminal profile as given in [i.2].
- System option map (16QAM, 32APSK, waveformFlex, lowerCarrier Switch, slotterAlohaTraffic …) that maps the optional features supported by the terminal for supervision following the nomenclature provided in [i.2].
- Features supported by the terminal (FeaturesMap field in the table below).
- Lower layer capabilities that are advertised during logon following the format of [i.3].
- Higher layer capabilities that are advertised during logon (to be reviewed against system features and system option map).
- Network topology support: star transparent, mesh regenerative, mesh transparent or hybrid, an indication of the network topology modes supported by the terminal.

- Transmission and reception encapsulation modes: GSE, RLE, ATM or MPEG. The encapsulation modes are given by the SNO, as a system configuration. Any terminal compatible with RCS2 should comply with RLE and GSE. The ATM and MPEG modes given are only set for backward compatibility purposes.

Table 11.4 proposes modifications to the System Configuration MIB group as described in [i.1].

**Table 11.4**

| Element | Range | Description | Access Rights | M&C Actor |
|---|---|---|---|---|
| SystemProfileMap | Cosumer(0),<br>SOHO(1),<br>Multi-dwelling (2),<br>Corporate (3),<br>SCADA (4),<br>Backhaul (5),<br>Institutional (6) | Indicates RCST supported profiles as bit map flags, where:<br>-0 not supported<br>-1 supported. | Installer (RO)<br>SNO (RO)<br>SVNO (RO) | This field should represent the profiles that can be supported by the terminal from factory.<br>The SNO/SVNO should be aware of the supported profiles.<br>The SNO RCST template for provisioning (that specifies the RCST profile) should be in accordance to the RCST supported profiles given in this field. |
| OptionMap | 16QAMrtn (0),<br>32APSKfwd (1),<br>waveformFlex (2),<br>contentionSync(12),<br>nomarclFec(13),<br>multiTs(14),<br>qsTs(15) | Minimum list of system options, given for supervision. | Installer (RO)<br>SNO (RO)<br>SVNO (RO) | Options provided from factory. SNO/SVNO should be aware of these values. |
| FeaturesMap | qpsk_8psk_cpmRtn (0),<br>refWaveforrms (1),<br>customWaveforms (2),<br>waveformBound (3),<br>waveformToTimeslot (4),<br>eirpPowerCtrol (5),<br>constantPowerCtrl (6),<br>fwdLinkDvbs2 (7),<br>fwdLinkSingleGS (8),<br>fwdLinkTSPacketStream (9),<br>fwdLinkMultipleStreams (10),<br>gseBBFrameCRC32 (11),<br>damaTraffic (12),<br>unsolicitedDATraffic (13),<br>slottedAlohaLogon (14),<br>recombinedDAMA (15),<br>raReplicas (16),<br>inbandSignalling (17),<br>signallingDAtimeslots (18),<br>SCPC (30),<br>space3 (31),<br>mobile (32) | These are the features supported by the terminal, given to the NCC/NMC for information. | Installer (RO)<br>SNO (RO)<br>SVNO (RO) | Features provided from factory. SNO/SVNO should be aware of these values. |

| Element | Range | Description | Access Rights | M&C Actor |
|---|---|---|---|---|
| LowerLayerCapabilities | multipleGS1(0),<br>multipleGS2(1),<br>reserved1(2),<br>fullRangeFLMODCOD (3),<br>fullRangeRLMODCOD(4),<br>fastCarrierSwitching (5),<br>carrierSwitchingClass1(6),<br>carrierSwitchingClass2(7),<br>EsN0powerCtrl(8),<br>constantPowerSpectrumDensity(9),<br>slottedAlohaTraffic(10),<br>crdsaTrafficSupport (11),<br>reserved2(12),<br>reserved3(13),<br>reserved4(14),<br>customCCCPMwaveform(15),<br>service1(16),<br>service2(17),<br>service3(18),<br>service4(19),<br>nbrofL2ifs(20),<br>nbrofL2ifs(21),<br>nbrofL2ifs(22),<br>nbrofL2ifs(23),<br>SWversion1(24),<br>SWversion1(25),<br>SWversion1(26),<br>SWversion1(27),<br>SWversion1(28),<br>SWversion1(29),<br>SWversion1(30),<br>SWversion1(31) | Lower layer capabilities following Table 8.5 from [i.3]. Each field is one flag. (bit). Information provided by the RCST to the NCC during logon. | Installer (RW)<br>SNO (RO)<br>SVNO (RO) | These flags should be set in accordance to the capabilities activated in the terminal<br>LL capabilities information is provided by the RCST to the NCC during logon.<br>This information is required by the NCC, to determine the operation parameters of the terminal.<br>The SNO should configure the terminal in accordance to these flags and the RCST provisioning. |
| HigherLayerCapabilities | ipv4ipv6Support (0),<br>multicastFwd (1),<br>enhMulticast(2)<br>dynamicMulticast (3),<br>diffservQoS (4),<br>mplsSupport (5),<br>snmpv2c(6),<br>snmpv3 (7),<br>dynamicConnectivity(8),<br>transecHooksSupport (9),<br>dynamicRouting (10),<br>ospfSupport (11),<br>firewall (12),<br>multiSVNO (13)<br>VLAN(14),<br>dhcpLAN (15),<br>motorControl (16),<br>sddp (17),<br>pepNegotiationProtocol (18),<br>authenticatedLogon (19),<br>mesh (20),<br>reserved (21),<br>reserved (22),<br>reserved (23) | Higher layer capabilities. Information provided by the RCST to the NCC during logon. | Installer (RW)<br>SNO (RO)<br>SVNO (RO) | These flags should be set in accordance to the capabilities activated in the terminal. These capabilities should be known by the SNO and the SVNO, to be able to adjust the HL operation accordingly. |

| Element | Range | Description | Access Rights | M&C Actor |
|---|---|---|---|---|
| PointingAlignmentSupport | 0 – Nominal CW EIRP in the pointing direction<br>1 – Supported pointing alignment methods<br>- (1) Burst probe, and CW probe by fixed non-configurable EIRP<br>- (2) Burst probe, and CW probe by configurable EIRP | New proposed 2 byte field that indicates the support of pointing alignment probing. Parameter is proposed to be moved to the installation group together with the rest of the alignment parameters. | Installer (RW)<br>SVNO (RO)<br>SNO (RO) | Flag used to inform the NCC the kind of alignment procedure supported by the RCST. The type of alignment is selected during the alignment process. |
| NetworkTopologySupport | starTransparent (0),<br>meshRegenerative (1),<br>meshTrasnparent (2),<br>hybrid (3) | Network topology read-only parameter | Installer (RO)<br>SNO (RO)<br>SVNO (RO) | Flags that indicate the type of topologies supported by the terminal. A change of topology may be linked to a new software version. SNO/SVNO should be aware of this value. |
| NetworkEncapsulationModeTx | ATM(1),<br>MPEG(2),<br>RLE(3),<br>GSE(4) | Encapsulation mode for transmission<br>If the terminal is RCS2 compliant, it should be able to support the 4 different possibilities. | Installer (RW)<br>SNO (RW)<br>SVNO (RO) | Value configured by the installer. It can then be reconfigured by the SNO. |
| NetworkEncapsulationModeRx | ATM(1),<br>MPEG(2),<br>RLE(3),<br>GSE(4) | Encapsulation mode for reception<br>If the terminal is RCS2 compliant, it should be able to support the 4 different possibilities. | Installer (RW)<br>SNO (RW)<br>SVNO (RO) | Value configured by the installer. It can then be reconfigured by the SNO. |

### 11.1.3 SNMP initial configuration

The RCST first installation should include a set of SNMP parameters to allow the RCST reception of SNMP commands from the primary NMC even if a successful logon into the network has not been yet performed.

This would require that:

- The RCST accepts SNMP commands through the satellite interface. The RCST may have several interfaces. SNMP access filters are applied to RCST IfIndex 1.
- The NMC sends SNMP commands using the Hardware ID (6 bytes) address that uniquely identifies the terminal in the network. Once the terminal has logon and has the SVN-MAC (3 bytes) for SVN0, the SNMP commands can use this SVN-MAC for SNMP traffic.
- The NMC SNMPv2c community is configured in the terminal through the configuration of snmpCommunityTable as defined in the "SNMP Community MIB Module" clause of [i.32] and the snmpTargetAddrTable is defined in the "Definitions" clause of [i.33].

The RCST may create one row in snmpTargetAddrTable for each SNMPv2c Transport Address Access.

SNMP access is controlled and specified by the MIB objects in [i.35] through [i.34], and [i.32].

## 11.2 RCST alignment

The RCST alignment process may include two different stages:

- forward link (FL) acquisition prior to enabling transmission on the return link;
- return link (RL) required only if FL pointing accuracy is achieved and to perform an initial MAC logon.

### 11.2.1 RCST forward link antenna alignment configuration

Following the description of the forward link antenna alignment in [i.1], this process will require the following parameters:

- MaxFwdAlignThrExcDuration: the duration of the time interval during which FL alignment accuracy should be achieved (part of the Installation MIB group).
- Max Fail: Maximum number of alignment failures (part of the Installation MIB group). The corresponding counter is incremented every time the state machine re-visits the FwdAlignment state.
- Fwd_link_snr_threshold (part of the Pointing Alignment Control Descriptor): the FL SNR threshold value to be reached to ensure FL successful alignment, value required for FL alignment accuracy. This parameter is proposed for inclusion in the Installation Group of the MIB.
- Alignment Population ID: A different population ID to be used during the alignment process. This will be provided by the NCC while negotiating the alignment parameters. Could be saved in the RCST MIB for supervision, as for now, this parameter is not included in the MIB.
- Start-up downlink TDM (administratively configured and selected by the RCST): The RCST should tune to the start-up in the operational TDM (Flink Configuration parameters). From there, the RCST can request the alignment process.

[i.4] proposes several suitable mechanisms to ensure forward link accuracy:

- manual procedure support by acoustic or visual feedback directly related to the power measurements of the received RF signal (CNR);
- automated procedure via motorized antenna as detailed in clause 10 of HLS [i.1].

The type of FL alignment mechanism is linked to the flag motorControl(16), part of HL capabilities. If activated the RCST will inform the NCC whether it has or not a motorized antenna.

To sum up these are the parameters that need to be set up by the installer to achieve forward link acquisition and alignment.

**Table 11.5**

| Functional Group | dvbRcs2Installation | | | | |
|---|---|---|---|---|---|
| **Element** | **Type** | **Range** | **Description** | **Access Rights** | **M&C Actor** |
| MaxFwdAlignThrExeDuration | Unsigned 32 | (0) notAcquired, (1) acquired | The duration of the time interval during which FL alignment accuracy should be achieved | Installer (RW) SNO/SVNO (RO) | SNO system parameter, that applies to all SNO's terminals. |
| MaxFail | Counter | | Maximum number of alignment failures allowed | Installer (RW) SNO/SVNO (RO) | SNO system parameter, that applies to all SNO's terminals. |
| **Functional Group** | **dvbRcs2SystemConfig** | | | | |
| HigherLayerCapabilities motorControl (16), | | 0 – manual 1 - motorized antenna | Whether the terminal has or not a motorized antenna | Installer (RW) SNO (RO) SVNO (RO) | These flags should be set in accordance to the capabilities activated in the terminal. These capabilities should be known by the SNO and the SVNO, to be able to adjust the HL operation accordingly. |
| OduRxBand | INTEGER | | LNB high band / Low band selector. High band corresponds to the emission of an 18-26 kHz tone with 0,4-0,8 Vpp in the Rx IFL cable. **Required for forward link acquisition.** | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision, only in RCST MIB |
| OduRxLO | INTEGER 32 | | ODU reception local oscillator frequency. **Required for forward link acquisition.** | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision, only in RCST MIB |
| OduTxLO | INTEGER 32 | | ODU transmission Frequency of Block Up-Converter Local Oscillator. **Required for forward link acquisition.** | Installer (RW) SNO/SVNO (RO) | SNO/SVNO access for supervision, only in RCST MIB |

The FL alignment is performed with the operational population ID, taking the first valid entry in the forward link configuration from the Flink Configuration group. After FL alignment, the RCST is able to filter all the necessary control information related to the RCS network and can request a further Return link alignment.

### 11.2.2 Return link alignment

After a successful forward link acquisition, the RCST is aware of the RCS2 network properties. At this point the terminal can start transmitting, and even require a return link alignment.

The RL alignment can be done in two different ways:

- Based on Installation Burst (IB)
- Based on Continuous Wave (CW) transmission

Both ways could be performed either automatically or manually.

In [i.1] the type of pointing alignment support is defined in the System Configuration MIB group as. [i.3] also defines this parameter as a logon element type (passed to the NCC in the logon request), as follows:

- Pointing Alignment support: the different ways that the RCST may support link alignment operations

**Table 11.6**

| MSB | LSB | Supported pointing alignment methods |
|---|---|---|
| 128-255 | User defined | User defined |
| 2-127 | Reserved | Reserved |
| 1 | Nominal CW EIRP in the pointing direction, in dBm | Burst probe, and CW probe by fixed non-configurable EIRP |
| 0 | Reserved | Burst probe, and CW probe by configurable EIRP |

However this table is incorrect. New proposed 2 byte field that indicates the support of pointing alignment probing:

0 – Nominal CW EIRP in the pointing direction

1 – Supported pointing alignment methods:

- (1) Burst probe, and CW probe by fixed non-configurable EIRP
- (2) Burst probe, and CW probe by configurable EIRP

The type of pointing alignment support by the RCST is configured by the installer, and should be reflected in the RCST MIB for supervision. For a better organization of parameters, this configuration should be placed under the installation group.

**Table 11.7**

| Functional Group | dvbRcs2SystemConfig | | | | |
|---|---|---|---|---|---|
| **Element** | **Type** | **Range** | **Description** | **Access Rights** | **M&C Actor** |
| dvbRcs2PointingAlignment Support | INTEGER32 | 0 – Nominal CW EIRP in the pointing direction (1 byte)<br>1 – Supported pointing alignment methods<br>- (1) Burst probe, and CW probe by fixed non-configurable EIRP<br>- (2) Burst probe, and CW probe by configurable EIRP | New proposed 2 byte field that indicates the different ways that the RCST may support link alignment operations. Parameter is proposed to be moved to the installation MIB group together with the rest of the alignment parameters. | Installer (RW) SNO/SVNO (RO) | Flag used to inform the NCC of the kind of alignment procedure supported by the RCST. The type of alignment is selected during the alignment process. |

To complete the RL alignment configuration, and to allow any RCST to do a successful logon, the RCST would need:

**Table 11.8**

| Functional Group | dvbRcs2SystemConfig | | | | |
|---|---|---|---|---|---|
| **Element** | **Type** | **Range** | **Description** | **Access Rights** | **M&C Actor** |
| sysLocation | DisplayString | | GPS position of the RCST ODU expressed as longitude, latitude and altitude. The string has 31 characters in the following format <xx.xxx>, <a>, <yyy.yyy>, <b>, <zzzz.z>, M, where x,y and z represent digits, a=N or S, b= E or W. | Installer (RW) SNO (RW) SVNO (RO) | SNO remote access for recovery |
| **Functional Group** | **dvbRcs2RtnConfiguration** | | | | |
| RtnConfigMaxEirp | Integer32 | | Maximum value of EIRP that the terminal can reach | Installer (RW) SNO (RW) SVNO (RO) | SNO remote access for recovery |
| RtnConfigDefIfLevel | Integer32 | | Starting power level for IF | Installer (RW) SNO (RW) SVNO (RO) | SNO remote access for recovery |

During logon, the RCST informs the NCC if it can support one or more of the RL alignment operations by means of the pointing alignment support indicator.

## 11.3 RCST logon and first commissioning

If no Pointing Alignment Support descriptor is present in the TIM-B, the RCST can proceed normally with FL acquisition and logon attempt.

After successful FL acquisition, the RCST should verify the status (dvbRcs2AlignmentStatus element of the State group) of earlier pointing alignment. If done, no alignment process is required and the RCST can continue with the logon.

During FL acquisition, the RCST receives the NIT, RMT, NCR, SPT, FCT2, BCT, TIM-B, TBTP2, having all the necessary control information related to the operation in the RCS network.

The RCST checks the Lowest Software Version descriptor matching its RCST HID. This information is included in the TIM-B. The RCST can only proceed with the MAC logon if its current operational SW version defined by implementation rules is considered sufficient.

The descriptor contains the following information:

- oui: indicates a group of RCSTs by reference to an OUI matching the OUI used in the RCST HID;
- swdl_mcast_address/port: identifies the IPv4 multicast address and UDP destination port for a SW download multicast service;
- sw_version: the field indicating the lowest SW version associated with the OUI.

The following set of parameters is reflected in the SDDP configuration group of the RCST MIB:

- Operational SW version
- MinSwVersion
- IP information for downloading an new SW version

- IPv4 address (of an IP multicast stream) and UDP port

- A flag parameter to indicate to the RCST whether or not to ignore the SW version notified in the TIM-B. This flag needs to be included in the SDDP group

- Additionally, there is a backup SWversion in the state group

**Table 11.9**

| Functional Group | dvbRcs2SDDPconfiguration | | | | |
|---|---|---|---|---|---|
| **Element** | **Type** | **Range** | **Description** | **Access Rights** | **M&C Actor** |
| SwVersion | Unsigned32 | | Current SW version in the SW distribution carousel, respective to the manufID and vendor specific parameters | RW | |
| MinSwVersion | Unsigned32 | | Indicates the minimum SW version required for log-on, as received in the Lowest Software Version descriptor (TIM-B) | RW | |
| MgroupType | InetAddressType | | | RW | |
| MgroupAddress | InetAddress | | | RW | |
| MgroupPrefixLength | InetAddressPrefix Length | | | RW | |
| Port | InetPort | | | RW | |
| **Functional Group** | **dvbRcs2State** | | | | |
| dvbRcs2RCSTAlternateSoftwareVersion | snmpAdminString | | Alternate (backup/new) RCST software version ([i.39]) | RO | |

If the current SW version is insufficient, the RCST cannot log on, but perform the necessary actions to automatically load or acquire another operational SW version. The HLS specification recommends the usage of SDDP to download the new SW version. The SVN mask used by the multicast stream dedicated to SW download can be located by the RCST through the MMT2 or the mapping method indicated in the Logon Response descriptor.

After successful check of the correct RCST SW version, the RCST is ready to start a logon procedure. The LL specification [i.3] introduces two variants of the logon procedure:

- basic logon

- logon at large timing uncertainty

The procedure and parameters required for the basic logon is analyzed hereafter.

The RCST sends a logon request in a logon timeslot, either using random access or a logon timeslot dedicated to the RCST. This request includes:

- indication of the type of logon (entry type = 0x1 binding user to HW and network, see Pointing Alignment Support descriptor in [i.3])

- indication of the network status of the RCST as it perceives it (LSB of access status is 1 indicating that NCC has confirmed pointing alignment, see Alignment Control Types in [i.3])

- RCST HID (concerns only random access)

- a field indicating the lower layer capabilities of the RCST

- in addition, a field indicating the higher layer capabilities of the RCST

The higher layer capabilities field should follow the format already detailed in clause 11.1.2.

For supervision, the last type of logon requested and the indication of the network status of the RCST should be reflected in the RCST MIB (state group). RCST HID may be included as part of the MIB, as a RO parameter, part of RCS2 System group (i.e. new element to be added).The last logon entry type should also be saved in the status RCST MIB group (need to be included).

**Table 11.10**

| Functional Group | dvbRcs2State | | | | |
|---|---|---|---|---|---|
| **Element** | **Type** | **Range** | **Description** | **Access Rights** | **M&C Actor** |
| typeOfLogon | INTEGER | Basic (0), LargeTiming (1) | Two variants of logon procedure exist, the basic procedure and a procedure extension called Logon at Large Timing. (RCS2) | Installer (RW) SNO (RW) SVNO (RO) | First value provided in the first installation. Could be changed by the SNO. |
| dvbRcs2AlignmentStatus | INTEGER | (0) not confirmed aligned, (1) confirmed aligned | RCST flag that reflects the alignment status given by the NCC during logon. RCS2 | Installer, SNO, SVNO (RO) | |
| dvbRcs2SubscriptionStatus | INTEGER | (0) NotConfirmedSubscription (1) ConfirmedSubscription | Flag to reflect the RCST subscription status given by the NCC at logon. (RCS2) | Installer, SNO, SVNO (RO) | The flag reflects the information provided during logon, it is saved in the MIB for supervision. |
| dvbRcs2HLSinitialization | INTEGER | (0) HL not Initialized (1) HL initialized by the SNO (2) HL initialized by the SVNO | HL should be initialized by the SNO during logon. The SVNO may afterwards modify/complete the HL configuration. For that it should change first this status to (0), and once finished change it to (2) | Installer (RO) SNO (RW) SVNO (RW) | This value reflects the status of HL configuration. Only the SNO or SVNO may modify this value. |
| dvbRcs2ComissionedStatus | INTEGER | (0) Not confirmed commissioned (1) NCC indicates the commissioning is completed | RCST commissioned status. The flag can be raise by loading a new configuration file. At a change of NIT or RMT, the RCST changes this flag to "Not confirmed commissioned" (RCS2) | Installer (RO) SNO (RW) SVNO (RW) | This flag is set during the logon phase. But in order to allow remote configuration by other means (not only L2S but SNMP or configuration file from the SNO/SVNO), these flags are RW. The SNO and SVNO should change these flags following the similar rules as in the logon. |

The NCC TIM-U response includes:

- logon response descriptor, initializing the RCST for normal operation in the network (see Table 6.11);
- control assign descriptor, indicating the MF-TDMA sync thresholds;
- correction message descriptor, indicating initial corrections in timing, frequency, and power relative to transmission of the logon request bursts;
- lower layer service descriptor, that initializes the LL services;
- a Network Layer Info descriptor (NLID) for additional information, by default provided in SNMP format;
- conditionally, a Higher Layers Initialization descriptor;
- optionally, a DHCP Option descriptor with the MTU for the return link, sent in TIM U or in TIM B.

### 11.3.1 Higher layers initialization

This clause describes what information is needed for the HL to be initialized.

As part of the logon response, the following fields are relevant for HL initialization:

- RCST_access_status: This status can be used by the NCC to signal that the RCST is not commissioned or has its Higher Layers not initialized.
  - Access status = 0011
    - LSB of access status is 1 indicating that NCC has confirmed pointing alignment
    - xx1x indicates that NCC confirms that the user is associated with the RCST (User ID indicated in the logon request)
    - x0xx indicates that the HL have not been initialized
    - 0xxx indicates that the NCC has not confirmed that the commissioning is complete
- Unicast RCS-MAC addresses/SVN Mask for higher layers. unicast_rcsmac_count indicates the number of unicast RCS-MAC addresses that are assigned to the RCST; For each of them the logon response contains a:
  - svn_prefix_size: A 5 bit field that indicates the number of most significant bits of the associated unicast RCS-MAC that holds the SVN number
  - unicast_rcsmac A 24 bit field that assigns one unicast RCS-MAC to the RCST. The SVN bits constitute a bit field that holds the SVN number of the RCS-MAC

The Higher Layers Initialization descriptor (if access status indicates that the HL are not initialized this descriptor is included) is used by the NCC to initialize each of the layer 2 RCST interfaces for IPv4 based M&C. This way the SNO, initializes each one of the RCST's SVN interfaces with a different IPv4 address, being, each one of them, an additional traffic interface.

- sat_l2if_count: indicates the number of layer 2 interfaces that are initialized, and for each of them the HLID contains:
  - rcs_mac: A 24 bit field that provides a reference to one satellite side layer 2 interface by its dedicated RCS-MAC address
  - l2if_ipv4_m&c_address: A 32 bit field that indicates the IPv4 M&C address associated to a satellite side layer 2 interface; Overrides the initial SNMP configuration
  - hl_offer_stream_ipv4_mcast_identification: A 32 bit filed that indicates the IPv4 multicast stream to be used to discover the higher layer support offer. Used for PEP advertisement
  - hl_offer_stream_port_number: A 16 bit field that indicates the port number used for indicating the higher layer support offer. Used for PEP advertisement
  - higher_layer_ pep_switch_off: A flag that when set to '1' indicates that the RCST should switch off all higher layer interception PEPs for the respective satellite side layer 2 interface and apply the native protocols unmodified. After successful logon, PEP negotiation will be used to establish the PEP type per RCST interface

The SVN_0 is the one used only for management from the SNO, and it should be the first entry of the loop.

The RCST should support at least one traffic interface. The minimum number of entries in this loop should be two, the first one linked to SVN_0 for management from the SNO, and a second entry associated to traffic. More entries can be added, corresponding to the additional SVNs. The field "l2if_ipv4_m&c_address" really corresponds to the IPv4 address for traffic, but in addition, it can be used for management from the SVNO. The SVNO is free to select which traffic SVN to use for management, although, most likely, the decision would depend on the type of traffic that each SVN is carrying.

The SVN-MAC and SVN mask allows an RCST to identify the corresponding SVN number. The svn_prefix_size provided in the logon response, indicates the number of most significant bits of the associated RCS-MAC that holds the SVN number.

The HLS specification also mentions that:

Within the OVN an RCST should be assigned one or more IPv4 address corresponding to the configured SVN-MAC labels. The IPv4 address should be unique within a VRF Group. In addition, the RCST should allow the SVN-MAC interface to be assigned an IPv6 address and may support other network addresses.

NOTE: An RCST that is assigned multiple SVN-MAC labels corresponding to multiple traffic SVNs will normally also be assigned a separate IP address for each SVN-MAC (e.g. an IPv4 or IPv6 address). These addresses may be presented on separate physical LAN interfaces or separate VLAN sub-interfaces providing connectivity to multiple routed networks.

Following [i.1] there should be one IPv4 or IPv6 address for traffic assigned per SVN-MAC label. Right now, IPv6 addresses are not considered in this descriptor. LL specification does not include any provision for IPv6 addresses as SVN interface address, even if it is required that the terminal should be capable of transmitting and receiving IPv6 traffic. This configuration could be solved by HLS new descriptors or other means of configuration.

Higher layer initialization description information should be persistent across RCST restart and reboot.

The Network Layer Info descriptor (if access status indicates that the HL are not initialized this descriptor may be included in the TIMU message) provides a mechanism by which network level information can be passed (transparently through the lower layers) to the Management Plane of the RCST during, or prior to, the start-up configuration phase of logon. The message body datagram will take the form of an SNMP message, and will be formatted according to [i.69] and [i.70], and the PDU type should be a SetRequestPDU.

To complete the HL configuration and according to the HL capabilities, the NCC will use the NLID. The minimum set of NLID parameters should cover:

- multicast mode: forwarding enabled/disabled, IGMP proxy, IGMP querier, MLD, etc., for each traffic interface.
- QoS default configuration: default HL service, default IP classification table entries.
- Default OSPF configuration for each VRF group.

Alternatively these parameters could be sent by configuration file (SDDP) or using the multicast stream specified in the HLID (for the moment this stream is only used for PEP negotiation). These methods could also be used if additional HL configuration is required.

After the terminal has been commissioned and its Higher Layers initialized, the SVNO could, at any given time, put the RCST "HL Maintenance" mode (i.e. set the HL initialized flag in the state MIB group to 0, will re-send the NLID with the new configuration, and when finished, set back the flag to HL initialized.

#### 11.3.1.1 NLID fields

##### 11.3.1.1.1 Multicast

In the HL capabilities, the RCST indicates whether it supports dynamic multicast or not:

- dynamic multicast not supported. Enabling multicast reception in the SVN1 (in this example, the LAN interface only supports IPv4) of an RCST will imply creating a new row in the dynamic table vrfGroupTable:
  - vrfGroupIndex: 1 (1st row)
  - vrfGroupSVNnumber: 1 (configured by SNO in the NMC)
  - vrfSVNMAClabel: SVN1 (Octet string)
  - vrfGroupIfInterface: 1 (LAN interface number)
  - vrfGroupSVNMAC: RCST SVN-MAC of the LAN interface
  - vrfSVNmtu: 1500
  - vrfGroupIfInterface: LAN interface number

- vrfOSPFrouting: Enabled
- vrfOSPFrouterAddressType: IPv4
- vrfOSPFrouterAddress: DR Address (IP RCST-GW)
- vrfOSPFrouterPrefix: DR IP prefix
- vrfMulticastMappingMethod: MMT2 method
- vrfMulticastFwd: Enabled
- vrfMulticastRtn: Disabled
- vrfIgmpVersion: IGMPv2
- vrfIgmpQuerierLAN: Enabled
- vrfIgmpProxy: Disabled
- vrfIgmpQuerierSAT: Disabled
- vrfIgmpForward: Disabled
- vrfPimSM: Disabled
- vrfMldQuerierLAN: Disabled
- vrfMldProxy: Disabled
- vrfMldQuerierSAT: Disabled
- vrfMldForward: Disabled
- vrfGroupStatusRow: createAndGo

- dynamic multicast supported:
  - flag vrfIgmpProxy needs to be enabled
  - vrfOSPFrouterAddress and vrfOSPFrouterPrefix can be left empty and values can be dynamically taken by OSPF

#### 11.3.1.1.2 QoS default configuration

There are three QoS table in the RCST MIB: IPClassTable, HLServiceTable, and LLserviceTable.

The NLID could configure, in the HL initialization phase, entries in the IPClassTable and in the HLServiceTable. It is important that the IPClassHLSAssociation value corresponding to the default (match-all) IP class entry matches one existing HLServiceIndex.

The entry in the IPClassTable is used to compile all types of IP traffic:

- IPClassTable
- IPClassEntry
- IPClassIndex: 1
- IPClassDscpLow: 0
- IPClassDscpHigh: 63
- IPClassDscpMarkValue: 0
- IPClassIPProtocol: 255 (match all)

- IPClassSrcInetAddressType: ipv4(1)
- IPClassIPSrcInetAddress: 0.0.0.0
- IPClassSrcInetAddressPrefixLength: 0
- IPClassDstInetAddressType: ipv4(1)
- IPClassIPDstInetAddress: 0.0.0.0
- IPClassIPDstInetAddressPrefixLength: 0
- IPClassSrcPortLow: 0
- IPClassSrcPortHigh: 65535
- IPClassDstPortLow: 0
- IPClassDstPortHigh: 65535
- IPClassVlanUserPri: -1 (selectivity is inactive)
- IPClassVLANID: -1 (match any VLAN identifier)
- IPClassHLSAssociation: 1
- IPClassAction: 1 (forward de packet)
- IPClassOutOctets: Read-only
- IPClassOutPkts: Read-only
- IPClassRowStatus: createAndGo (the new row will become active after creation)

Here follows an example of an HLServiceTable entry (the parameters followed by a question mark are for the SNO to decide):

- HLServiceTable
- HLServiceEntry
- HLServiceIndex: 1
- HLserviceLLServiceAssociation: 1 (should be coherent with the LLServiceTable index)
- HLservicediffPolicyPHBindex: 0 (default PHB)
- HLservicePHBname: Default
- HLservicePriority: 0
- HLserviceMinRate: 0 Kbps
- HLserviceMaxRate: 2 000 Kbps
- HLserviceMaxIngressBurst 4000
- HLlserviceMinIngressBurst 20
- HLserviceMaxEgressBurst 4000
- HLserviceMaxDelay 30 sec.
- HLserviceQueueType: FIFO (0)
- HLserviceL3IfNumber: 1 (RCST LAN/VLAN interface)

- MaxLatency: 5 sec.
- LinkRetransmissionAllowed: packet retransmission not allowed (0)
- HLServiceRowStatus: createAndGo (the new row will become active after creation)

### 11.3.2 RCST commissioning

The access status of the RCST for a first logon after antenna alignment would be '0001'. Not till the higher layers have been initialized, and minimum RCST configuration is set, the RCST should change its status to HL initialize and commissioned, i.e. access status = 1111. This status is firstly set by the NCC during the logon.

The RCST commissioning and configuration is normally done during installation by RCST configuration file and is completed during logon thanks to the information provided in the TIM-U logon response message. However, if the commissioned-ok flag is not set, the RCST may block network forwarding of user traffic to/from the LAN interface. This allows further IP configuration. The RCST completes the configuration by enabling traffic forwarding when the commissioned-ok flag is set (e.g. by loading a new configuration or direct action to raise the flag).

The RCST can indicate that the status is "confirmed-commissioned" to the NCC if that NCC has previously indicated that the RCST has been commissioned (e.g. in a restart scenario), and the RCST has not, in the meanwhile, been re-commissioned towards another system or it has lost the previous alignment. If any of the latter occurs, the RCST should indicate that it is "not confirmed commissioned" in the logon request sent to the NCC. This allows the NCC to consider commissioning before allowing the RCST into the network.

The RCST commissioning status is reflected in the MIB state group. The status can remotely be checked by the SNO or SVNO.

After a successful commissioning status, the SVNO could decide to change the configuration of the terminal. For this, the SVNO should first change the commissioning status and then update the configuration of the terminal (e.g. by means of SNMP commands or new configuration file).

### 11.3.3 Logon and commissioning example

More details on the logon procedure are provided in Figure 11.1.

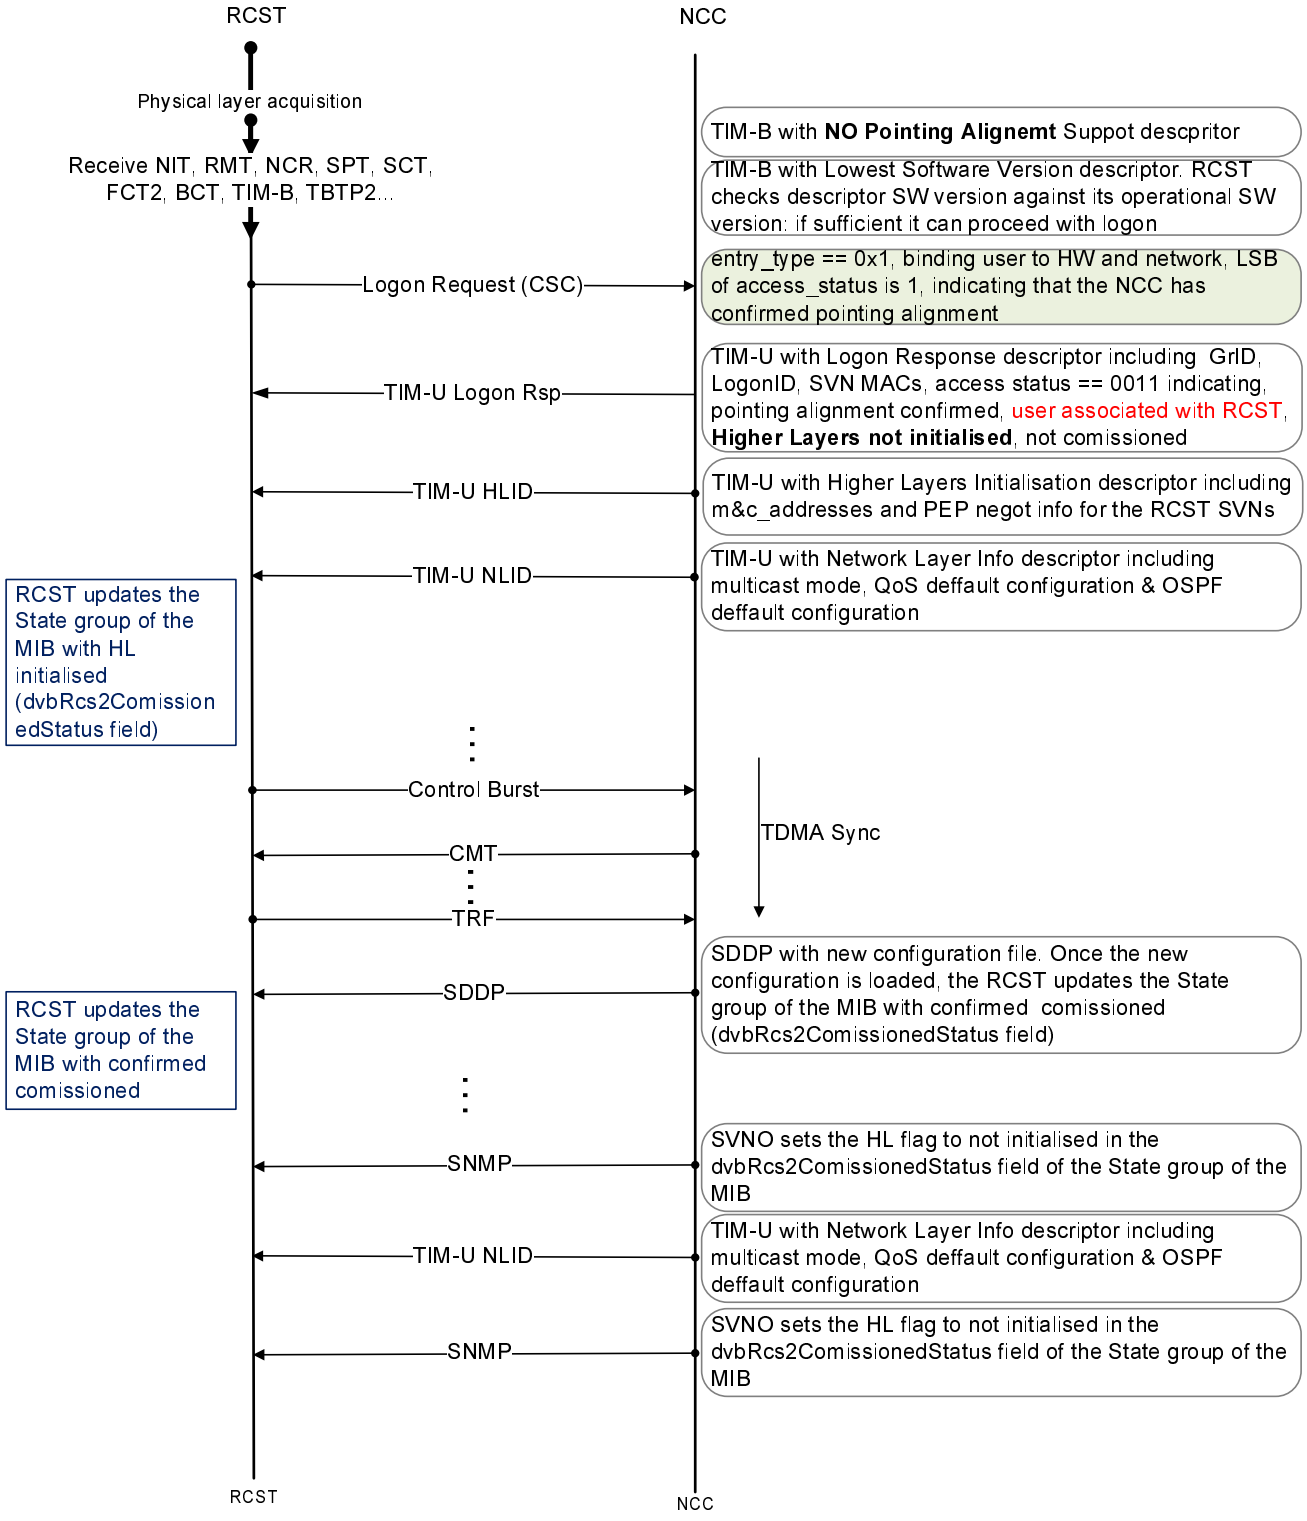


**Figure 11.1: Logon and commissioning sequence**

# 12 OSS-NMC interface and performance management guidelines

Network operators are deploying a range of different sub-networks to meet different demands in the telecommunications market (e.g. a combination of fixed broadband networks and satcom networks to provide internet to both urban and remote locations including the maritime segment). At the same time, sub-networks serving more or less the same purpose are gradually replacing each other over time, still living in parallel for some time (e.g. GSM, WCDMA, and LTE networks serving mobile communication). On top of this, especially in more dynamic, new markets, operators are growing by acquiring competitor networks, thus adding sub-networks of the same technology, but from different vendors to its operations.

In order to provide high-quality service at reasonable costs, operators will continuously aim to streamline an efficient and effective network operations organization. These will typically be organized with a Call Center (Level 1), 1st Line Technical Support (level 2), and Specialist Technical Support (Level 3).

Ideally, these organizational units would be the same for all operated sub-networks. However, due to required skill-levels (technology- and tool-wise), the operator often runs parallel organizational units doing the same job on different sub-networks/technologies.

This clause presents a methodology for standardized integration between OSS and the Network Management Center (NMC) of the satellite-based access network. The methodology makes use of existing 3GPP specifications. This may enable the re-use of the OSS applications that are already aligned with 3GPP in the terrestrial mobile networks.

## 12.1 OSS applications in mobile network operations

OSS applications used for operating mobile networks (e.g. GSM, WCDMA, and LTE) closely follow the PLMN (Public Land Mobile Network) management architecture, which is defined by the 3GPP (see [i.40] and [i.41]). The 3GPP PLMN management architecture is based on ITU-T TMN (Telecommunications Management Network standard from the ITU-T) which again can be seen as a refinement of the ISO FCAPS model.

The five management functions of FCAPS (**F**ault, **C**onfiguration, **A**ccounting, **P**erformance, and **S**ecurity) are still visible among the list of management functions of the 3GPP PLMN management architecture:

- Performance management
- Roaming management
- Fraud management
- Fault management
- Security management
- Software management
- Configuration management
- Accounting management
- Subscription management
- Quality of Service (QoS) management

A PLMN is often composed of equipment from a range of vendors. In order for integration to be successful, 3GPP proposes the use of Integration Reference Points (IRP) between Network Elements (NEs) and management functions. Figure 12.1 (from [i.41]) shows how Elements Manager (EM) and Network Manager (NM) should implement the IRP.

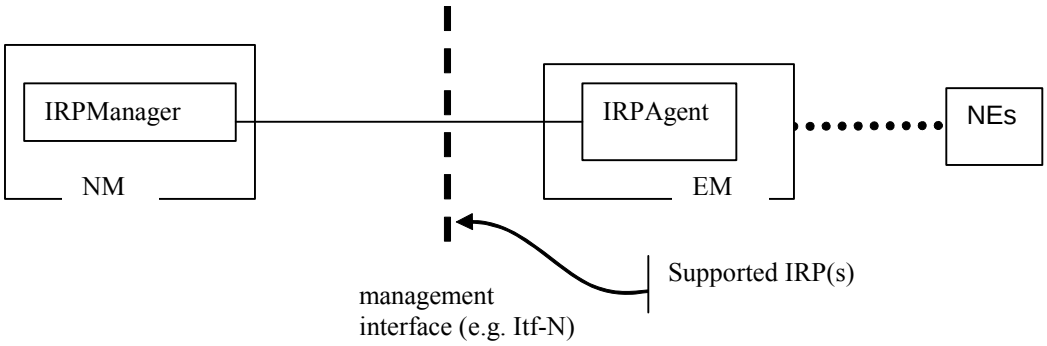


**Figure 12.1: 3GPP IRPs used in network element management**

XML is commonly used to transfer measurement results from Network Elements to the OSS as part of the Performance Management (PM). 3GPP has specified a PM XML file format in [i.42].

The PLMN management architecture can be re-used to a large extent in network operations with DVB-RCS2 satellite access.

## 12.2 Performance management concept

Performance Management (PM) aims to evaluate network behaviour. The physical and logical states of Network Elements (NEs) are measured and reported in data collected by the Element Manager (EM) function. This may be done according to some pre-defined time schedule.

Measurement data should be generated by NEs to meet the following purposes:

- measure the amount of user data and signalling traffic;
- verify the network configuration;
- measure the Quality of Service perceived by the user (e.g. throughput, round-trip-time, set-up time, etc.);
- measure resource availability and access control.

### 12.2.1 Measurement jobs

Measurement jobs executed in NEs are defined in the EM function. The definition includes scheduling the timing/frequency of measurement job execution, which specific data to measure/collect, and which (sub) components of the NE the measurement is valid for.

It should be possible to manage the measurement jobs in the EM. This entails the ability to start/stop/suspend/resume measurement jobs, and to view measurement jobs and their current status.

It should be possible to practically manage easily the many different measurement jobs in the network. This includes the ability to schedule the same measurement job "for all" NEs of a certain group or category (e.g. define the same measurement job to take place in all RCSTs).

### 12.2.2 Measurement results generation and storage

Each measurement job produces a number of results. The results should be contained in a measurement report associated with the measurement job.

Measurement result data needs to be kept in local storage in the NE or EM until it has been received by the NMC and OSS. Storage capacity and the duration for which data will be available locally at the NE or EM is implementation and operator dependent.

### 12.2.3 Measurement results transfer

Measurement results are transferred from the NE to the EM and NMC for storage, post-processing, and presentation in the OSS. There may be more than one OSS monitoring the same network, and serving multiple Satellite Network Operators (SNOs) and/or Satellite Virtual Network Operators (SVNOs). Therefore, results may need to be transferred to multiple destinations.

Measurement reports may be transferred from the NE to the EM in one of two ways:

a) notification-based transfer of reports when these are available,

b) on-demand transfer of reports when the EM (periodically) request these.

The measurement reports should be transferred from EMs to the NMC via bulk file transfer.

The NMC may store the files for a specified period of time (e.g. one hour) where it is available for the OSSs to collect them. Alternatively, the NMC may keep track of each destination OSS, notify these when report files are available, and remove these once all OSSs have notified the NMC that the reports have been processed.

### 12.2.4 Measurement report XML file format

Measurement report files may be stored in a well-defined XML file format aligned with 3GPP.

#### 12.2.4.1 3GPP XML file format

Table 12.1 shows XML tags specified by 3GPP.

**Table 12.1: XML tags used for performance measurement report example**

| XML tag | Description |
|---|---|
| measCollecFile | |
| fileHeader | |
| measData | |
| fileFooter | |
| fileHeader fileFormatVersion | |
| fileHeader dnPrefix<br>and<br>fileSender localDn | For the XML schema based XML format, the DN is split into the DN prefix and the Local DN (LDN) (see [i.43]). XML attribute specification "dnPrefix" may be absent in case the DN prefix is not configured in the sender. XML attribute specification "localDn" may be absent in case the LDN is not configured in the sender. |
| fileSender elementType | For the XML schema based XML format, XML attribute specification "elementType" may be absent in case the "senderType" is not configured in the sender. |
| fileHeader vendorName | For the XML schema based XML format, XML attribute specification "vendorName" may be absent in case the "vendorName" is not configured in the sender. |
| measCollec beginTime | |
| managedElement | |
| managedElement userLabel | For the XML schema based XML format, XML attribute specification "userLabel" may be absent in case the "nEUserName" is not configured in the CM applications. |
| fileHeader dnPrefix<br>and<br>managedElement localDn | For the XML schema based XML format, the DN is split into the DN prefix and the Local DN (LDN) (see [i.43]). XML attribute specification "localDn" may be absent in case the LDN is not configured in the CM applications. |
| managedElement swVersion | For the XML schema based XML format, XML attribute specification "swVersion" may be absent in case the "nESoftwareVersion" is not configured in the CM applications. |
| measInfo | |
| measInfoId | |
| granPeriod endTime | |
| job jobId | |
| granPeriod duration | For the XML schema based XML format, the value of XML attribute specification "duration" should use the truncated representation "PT*n*S". |
| repPeriod duration | For the XML schema based XML format, the value of XML attribute specification "duration" should use the truncated representation "PT*n*S". |
| measTypes<br>or<br>measType | For the XML schema based XML format, depending on sender's choice for optional positioning presence, either XML element "measTypes" or XML elements "measType" will be used. |
| measValue | |
| measValue measObjLdn | |
| measResults<br>or<br>r | For the XML schema based XML format, depending on sender's choice for optional positioning presence, either XML element "measResults" or XML elements "r" will be used. |
| suspect | |
| measCollec endTime | |
| measType p | An optional positioning XML attribute specification of XML element "measType" (XML schema based), used to identify a measurement type for the purpose of correlation to a result. The value of this XML attribute specification is expected to be a non-zero, non-negative integer value that is unique for each instance of XML element "measType" that is contained within the measurement data collection file. |
| r p | An optional positioning XML attribute specification of XML element "r", used to correlate a result to a measurement type. The value of this XML attribute specification should match the value of XML attribute specification "p" of the corresponding XML element "measType" (XML schema based). |

#### 12.2.4.2 Schema for performance measurement XML file format

XML schema, measCollec.xsd, specified in [i.42] may be used for this purpose.

#### 12.2.4.3 Example measurement report file in XML format

The following shows an example measurement report file in the XML format:

```
<?xml version="1.0" encoding="UTF-8"?>
<?xml-stylesheet type="text/xsl" href="MeasDataCollection.xsl"?>
<measCollecFile xmlns="http://www.3gpp.org/ftp/specs/archive/32_series/32.435#measCollec"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://www.3gpp.org/ftp/specs/archive/32_series/32.435#measCollec
http://www.3gpp.org/ftp/specs/archive/32_series/32.435#measCollec">
    <fileHeader fileFormatVersion="32.435 V7.0" vendorName="Company NN"
dnPrefix="DC=a1.companyNN.com,SubNetwork=1,IRPAgent=1">
        <fileSender localDn="SubNetwork=CountryNN,MeContext=MEC-Gbg-1,ManagedElement=RNC-Gbg-1"
elementType="RNC"/>
        <measCollec beginTime="2000-03-01T14:00:00+02:00"/>
    </fileHeader>
    <measData>
        <managedElement localDn="SubNetwork=CountryNN,MeContext=MEC-Gbg-1,ManagedElement=RNC-Gbg-
1" userLabel="RNC Telecomville"/>
        <measInfo>
            <job jobId="1231"/>
            <granPeriod duration="PT900S" endTime="2000-03-01T14:14:30+02:00"/>
            <repPeriod duration="PT1800S"/>
            <measTypes>attTCHSeizures succTCHSeizures
attImmediateAssignProcssuccImmediateAssignProcs</measTypes>
            <measValue measObjLdn="RncFunction=RF-1,UtranCell=Gbg-997">
                <measResults>234 345 567 789</measResults>
            </measValue>
            <measValue measObjLdn="RncFunction=RF-1,UtranCell=Gbg-998">
                <measResults>890 901 123 234</measResults>
            </measValue>
            <measValue measObjLdn="RncFunction=RF-1,UtranCell=Gbg-999">
                <measResults>456 567 678 789</measResults>
                <suspect>true</suspect>
            </measValue>
        </measInfo>
    </measData>
    <fileFooter>
        <measCollec endTime="2000-03-01T14:15:00+02:00"/>
    </fileFooter>
</measCollecFile>
```

## 12.3 Recommendations for DVB-RCS2 performance measurements

Performance measurements should meet the purpose of enabling the operator (SNO or SVNO) identify network elements with degraded performance as well as the root cause for degraded performances.

In summary, performance measurements should meet these requirements:

- Enable operator staff to discover degraded performance:
    - Service accessibility.
    - Service retainability.

  - Quality of service.

- Enable operator staff to isolate the root cause of degraded performance.
- Re-use existing DVB-RCS specifications (in particular, the DVB RCS2 MIB).
- Align with principles and specifications published by other telecommunications fora.

### 12.3.1 Performance measurements

Performance measurements presented here adhere to the principles in 3GPP performance measurement specifications (e.g. [i.44]).

Each measurement monitors/measures/counts a certain aspect of performance and has an obvious relation to other measurements such that comparable Key Performance Indicators (KPIs) can be defined.

Typically, it is desirable to scan the network on "success rates" on certain procedures in the network (e.g. the RCST logon procedure) in order to determine poorly performing network elements. Thus, it is necessary to count both the number of attempts and successes of the procedure. In order to isolate the root cause, it is also desirable to count the number of "failures" due to different causes when known.

For performance measurements monitoring a signal, it is desirable to measure both the transmitted signal strength, the received signal strength, and the quality/accuracy of the received signal in order to determine high losses or interference.

### 12.3.2 Impact on DVB-RCS2

This clause presents a minimum set of performance measurements, which may be extended by additional vendor-specific measurements. All performance measurements should be made available to the NMC/OSS through the mechanisms defined in clause 12.1.

## 12.4 Recommended performance measurements for DVB-RCS

In the following, measurements are defined for RCSTs in a DVB-RCS2-based satellite communications network.

### 12.4.1 Managed object classes

Although all measurements are done at RCST level, some will be done remotely at the RCST whereas others can be done at the central hub (NCC). The overall goal is to provide a sufficient set of measurements that enables efficient network monitoring where operations staff can easily compare measurements for all RCSTs across the network, no matter the equipment provider.

To separate remote and central measurements, the following assumes two different managed object classes, both representing RCSTs:

- RCSTRemote
- RCSTHub

Each corresponds to an Element Manager (EM) representing different aspects of the RCST Network Element (NE).

### 12.4.2 Measurement specification format

All measurements in the following are presented in the following structure:

- Textual description
- Collection Method (CC=Cumulative Counter, GAUGE, DER=Discrete Event Registration, SI=Status Inspection)
- Condition: The specific details/events causing an update to the measurement result

- Measurement units (e.g. seconds)
- Measurement identifier (as used in measurement result files)
- Managed object class (e.g. RCST)
- Technology generation (e.g. RCS2)

## 12.4.3 RCST accessibility

### 12.4.3.1 Number of Attempted Logons

- This measurement provides the number of attempted logons using DA and RA
- CC
- Receipt of a Logon burst (DA or RA) by the NCC from the RCST
- Each measurement is an integer value.
- RCST.AttLogon.DA
  RCST.AttLogon.RA
- RCSTHub
- RCS2

### 12.4.3.2 Number of Rejected Logons

- This measurement provides the number of rejected logons for an RCST for different causes (RESOURCE = no resource, ACCOUNT = account is valid or paid, OTHER)
- CC
- Transmission of a TIM-U message indicating rejected logon by the NCC to the RCST
- Each measurement is an integer value.
- RCST.RejLogon.RESOURCE
  RCST.RejLogon.ACCOUNT
  RCST.RejLogon.OTHER
- RCSTHub
- RCS2

### 12.4.3.3 Number of Acknowledged Logons

a) This measurement provides the number of acknowledged logon attempts

a) CC

b) Transmission of a TIM-U message indicating acknowledged logon by the NCC to the RCST

c) Each measurement is an integer value

d) RCST.AckLogon

e) RCSTHub

f) RCS2

#### 12.4.3.4 Number of Successful Logons

a) This measurement provides the number of successful logon attempts

b) CC

c) Receipt of a control burst message following logon acknowledgement by the NCC from the RCST

d) Each measurement is an integer value

e) RCST.SucLogon

f) RCSTHub

g) RCS2

#### 12.4.3.5 Number of Failed Logons

a) This measurement provides the number of failed logon attempts

b) CC

c) *Either:* No receipt of a control burst message following logon acknowledgement by the NCC from the RCST, *Or:* Receipt of another logon burst (DA or RA) by the NCC from the RCST, indicating that the RCST has not received a TIM-U message with logon acknowledgement

d) Each measurement is an integer value

e) RCST.FailLogon

f) RCSTHub

g) RCS2

#### 12.4.3.6 Number of Logoffs

a) This measurement provides the number of logoffs of different logoff causes specified in [i.3]

b) CC

c) Logoff message sent to RCST by NCC or autonomous silent logoff as per the logoff procedure described

d) Each measurement is an integer value

e) RCST.Logoff.NCC
RCST.Logoff.USER
RCST.Logoff.AUTO
RCST.Logoff.STANDBY
RCST.Logoff.SYNC
RCST.Logoff.FREQ
RCST.Logoff.INTERNAL
RCST.Logoff.OTHER

f) RCSTHub

g) RCS2

#### 12.4.3.7 Forward Link Bit Error Rate

a) This measurement provides the RCST Bit Error Rate (BER) of the Forward Link

b) SI

c) The average BER of the Forward link at the RCST within the granularity period

d) The result is an integer (0..63) where the meaning is:
0: BER = 0
1: -infinity < Log10(BER) < -6.1
2: -6.1 <= Log10(BER) < -6.0
...
61: -0.3 <= Log10(BER) < -0.2
62: -0.2 <= Log10(BER) < -0.1
63: -0.1 <= Log10(BER)

e) RCST.FwdBER

f) RCSTRemote

g) RCS2

#### 12.4.3.8 Forward Link Carrier-to-Noise Ratio

a) This measurement provides the RCST Carrier-to-Noise Ratio (CNR) of the Forward Link

b) SI

c) The average CNR of the Forward link at the RCST in 0,1 dB units within the granularity period

d) 0,1 dB

e) RCST.FwdCNR

f) RCSTRemote

g) RCS2

#### 12.4.3.9 Forward Link Received Power

a) This measurement provides the Forward Link Rx Power in the RCST

b) SI

c) The average RX Power of the Forward link at the RCST in 0,1 dBm units within the granularity period

d) 0,1 dBm

e) RCST.FwdRxPower

f) RCSTRemote

g) RCS2

#### 12.4.3.10 Return Link Received EbN0

a) This measurement provides the Return Link EbN0 of the RCST measured in the hub

b) SI

c) The average EbN0 of the Return link of the RCST in 0,1 dB units within the granularity period

d) 0,1 dB

e) RCST.RtnEbN0

f) RCSTHub

g) RCS2

#### 12.4.3.11 Return Link Transmitted EIRP

a) This measurement provides the Return Link EIRP of the RCST

b) SI

c) The average EIRP of the Return link in the RCST in dBW within the granularity period

d) dBW

e) RCST.RtnEIRP

f) RCSTRemote

g) RCS2

#### 12.4.3.12 Number of Capacity Requests

a) This measurement provides the number of solicited capacity requests sent by the RCST of the different capacity categories

b) CC

c) Receipt of capacity request by the NCC from the RCST

d) Each measurement is an integer value

e) RCST.CapacityRequests.VBDC
RCST.CapacityRequests.RBDC
RCST.CapacityRequests.AVBDC

f) RCSTHub

g) RCS2

#### 12.4.3.13 Number of Rejected VBDC Capacity Requests

a) This measurement provides the number of rejected VBDC capacity requests of different causes

b) CC

c) Evaluation by the NCC of a VBDC capacity request from the RCST where the capacity request is not met due to the following causes:

d) The capacity request backlog is full

e) The capacity request has expired

f) Resources are not available to satisfy the capacity request

g) Other reason

h) Each measurement is an integer value

i) RCST.VBDCCapacityRequestsFail.BACKLOG
RCST.VBDCCapacityRequestsFail.EXPIRED
RCST.VBDCCapacityRequestsFail.RESOURCE
RCST.VBDCCapacityRequestsFail.OTHER

j) RCSTHub

k) RCS2

#### 12.4.3.14 Number of Rejected RBDC Capacity Requests

a) This measurement provides the number of rejected RBDC capacity requests of different causes

b) CC

c) Evaluation by the NCC of a RBDC capacity request from the RCST where the capacity request is not met due to the following causes:

d) The capacity request backlog is full

e) The capacity request has expired

f) Resources are not available to satisfy the capacity request

g) Other reason

h) Each measurement is an integer value

i) RCST.RBDCCapacityRequestsFail.BACKLOG
RCST.RBDCCapacityRequestsFail.EXPIRED
RCST.RBDCCapacityRequestsFail.RESOURCE
RCST.RBDCCapacityRequestsFail.OTHER

j) RCSTHub

k) RCS2

#### 12.4.3.15 Number of Rejected AVBDC Capacity Requests

a) This measurement provides the number of rejected AVBDC capacity requests of different causes

b) CC

c) Evaluation by the NCC of a AVBDC capacity request from the RCST where the capacity request is not met due to the following causes:

d) The capacity request backlog is full

e) The capacity request has expired

f) Resources are not available to satisfy the capacity request

g) Other reason

h) Each measurement is an integer value

i) RCST.AVBDCCapacityRequestsFail.BACKLOG
RCST.AVBDCCapacityRequestsFail.EXPIRED
RCST.AVBDCCapacityRequestsFail.RESOURCE
RCST.AVBDCCapacityRequestsFail.OTHER

j) RCSTHub

k) RCS2

#### 12.4.3.16 Return Link Throughput

a) This measurement provides the total return link throughput of different capacity categories

b) CC

c) Sum of the throughput on all channels (all timeslots) in granularity period measured at the NCC

d) Kilobit

e) RCST.RtnThroughout

f) RCSTHub

g) RCS2

#### 12.4.3.17 Return Link Allocated Throughput

a) This measurement provides the return link throughput at different capacity categories

b) CC

c) Sum of allocated throughput on all channels in granularity period for the different capacity categories

d) Kilobit

e) RCST.RtnAllocThroughput.CRA
RCST.RtnAllocThroughput.FCA
RCST.RtnAllocThroughput.VBDC
RCST.RtnAllocThroughput.RBDC
RCST.RtnAllocThroughput.AVBDC

f) RCSTHub

g) RCS2

#### 12.4.3.18 Return Link Unused CRA Capacity

a) This measurement provides the return link unused CRA capacity

b) CC

c) Unused, available CRA capacity in the granularity period

d) Kilobit

e) RCST.RtnCRACapacityUnused

f) RCSTHub

g) RCS2

# 13 Dynamic connectivity protocol guidelines for mesh regenerative systems

Dynamic connectivity is supported in RCS2 thanks to the Dynamic Connectivity Protocol (DCP) as specified in Annex E of [i.1]. DCP is a control signalling protocol between the NCC and the RCST. This protocol is used when IP connectivity with the NCC is achieved after RCST logon and allows the mapping of IP parameters and policies to L2 parameters, and to dynamically set one or several mesh links within connectivity channels to an RCST according to set of values configured by L2S or management.

Mesh RCSTs (transparent or regenerative) support DCP protocol for mesh link establishment for DVB-RCS2 in Mesh Regenerative systems and Mesh overlay systems.

This clause introduces some recommendations on the usage of DCP over Mesh Regenerative Systems.

The Mesh System Descriptor (with tag 0xE1) is provided by the NCC in the TIM-B message. This descriptor indicates:

- whether or not the system is ready to process dynamic connectivity logon requests, and
- a list of frames that may be used for mesh traffic, for each superframe used for mesh.

Note that if the descriptor length is '0', then all frames can be used for mesh traffic.

The NCC may assume that the listed set of frames constitutes the RPLS for all the mesh receivers that are using the superframe.

# 13.1 DCP messages

The minimum set of messages to implement a valid DCP in a Mesh regenerative system is:

- Link Service Establishment Request by RCST
- Link Service Establishment Response by NCC
- Link Service Establishment Request by NCC
- Link Service Establishment Response by RCST
- Link Service Release Request
- Link Service Release Response
- Acknowledgement

The rest of the messages are optional, and are described in the [i.1] document.

## 13.1.1 DCP logon

The DCP Logon procedure can be started when IP connectivity with the NCC is achieved. It permits the mapping of IP parameters and policies to L2 parameters. It also allows to dynamically set connectivity channels to an RCST according to the set of values configured by management and L2S.

### 13.1.1.1 RCST DCP successful logon

A successful DCP Logon is achieved when the RCST receives confirmation by the NCC. The procedure is illustrated in Figure 13.1.

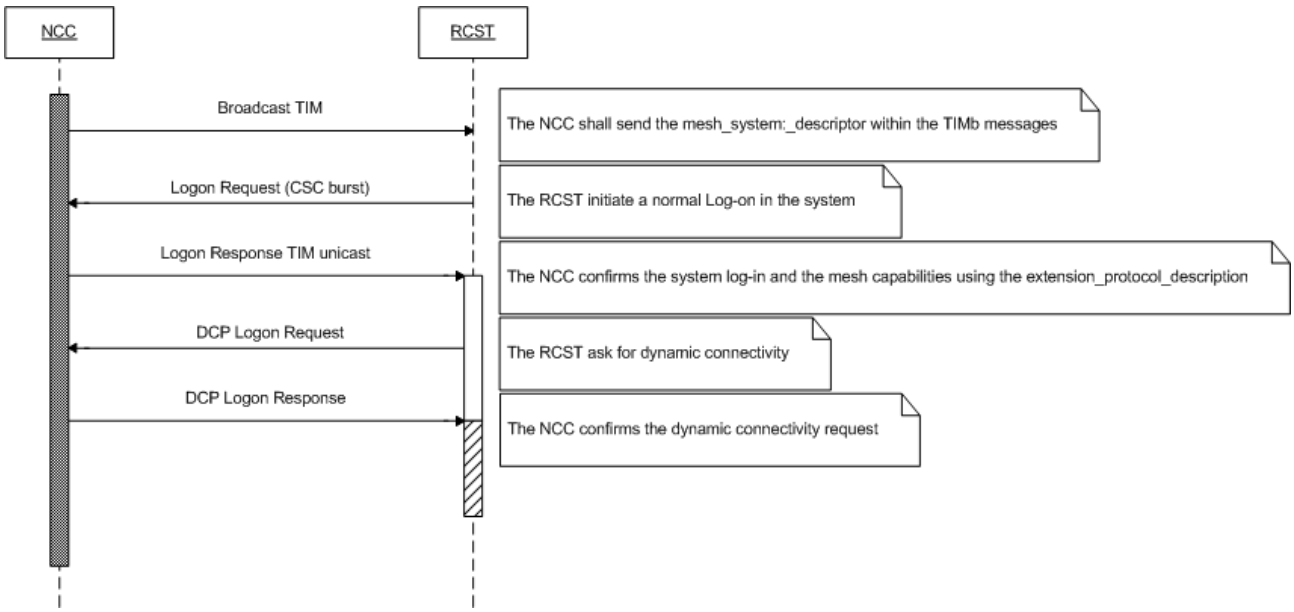


**Figure 13.1: Successful logon**

#### 13.1.1.2 RCST DCP failed logon

DCP logon may fail due to rejection from the NCC or due to DCP message loss(es). This is illustrated in Figure 13.2.

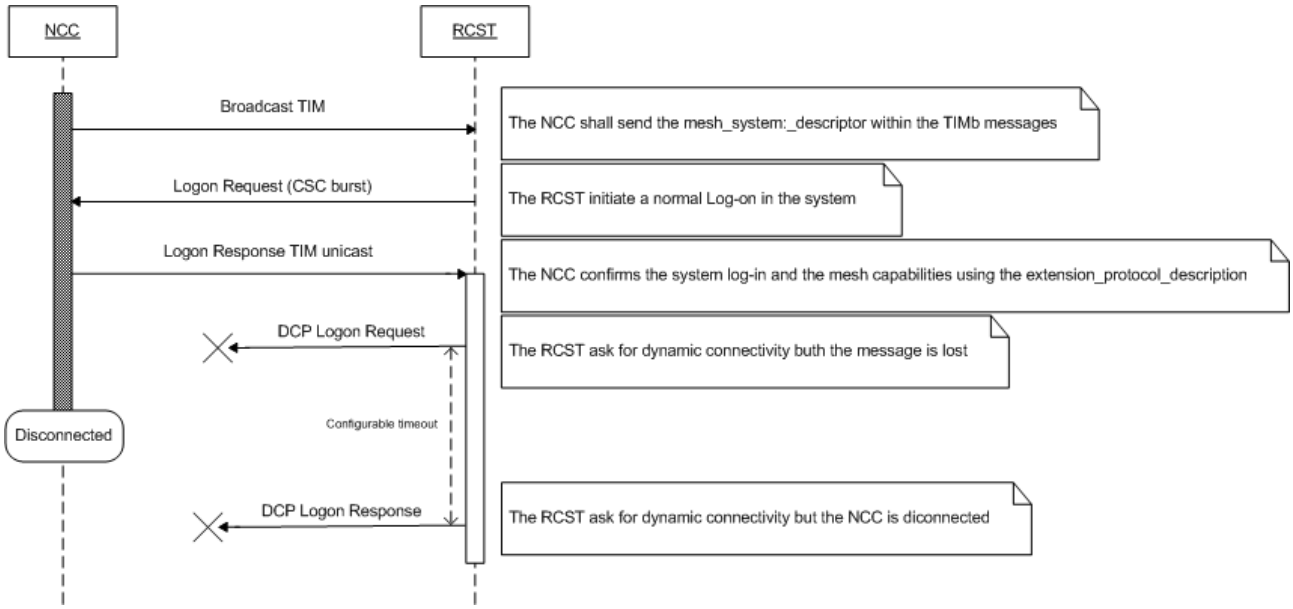


**Figure 13.2: Failed logon**

### 13.1.2 RCST DCP connections procedures

The basic dynamic connectivity procedures supported by a regenerative mesh RCST include:

- RCST-initiated bidirectional connections
- RCST-initiated unidirectional multicast connections
- NCC-initiated bidirectional connections
- NCC-initiated unidirectional multicast connections

RCST sends a LINK SERVICE ESTABLISHMENT REQUEST message upon receiving at its LAN interface an IP packet that cannot be mapped to an existing connection. More specifically, two triggers are listed below identifying the conditions under which the LINK SERVICE ESTABLISHMENT REQUEST messages is sent:

- addressing/routing trigger, if the packet matches an existing flow type (with defined IP CoS/PHB), but its next hop IP destination address does not match any of the existing connections;
- QoS trigger, if the packet's IP CoS/PHB does not match the service used in an existing connection using the same next hop IP address.

A packet can be forwarded to an active connection only when it is addressed to the same destination RCST and if its associated LL service matches that of the active connection.

The RCST sends a LINK SERVICE RELEASE REQUEST message for those active connections not carrying traffic in either direction after a configurable timeout.

Upon receiving a LINK SERVICE RELEASE REQUEST, an RCST sends a LINK SERVICE RELEASE RESPONSE to the peer RCST as acknowledgement of the request.

### 13.1.2.1 RCST DCP successful unicast connection

A successful unicast connection establishment procedure is described in Figure 13.3.

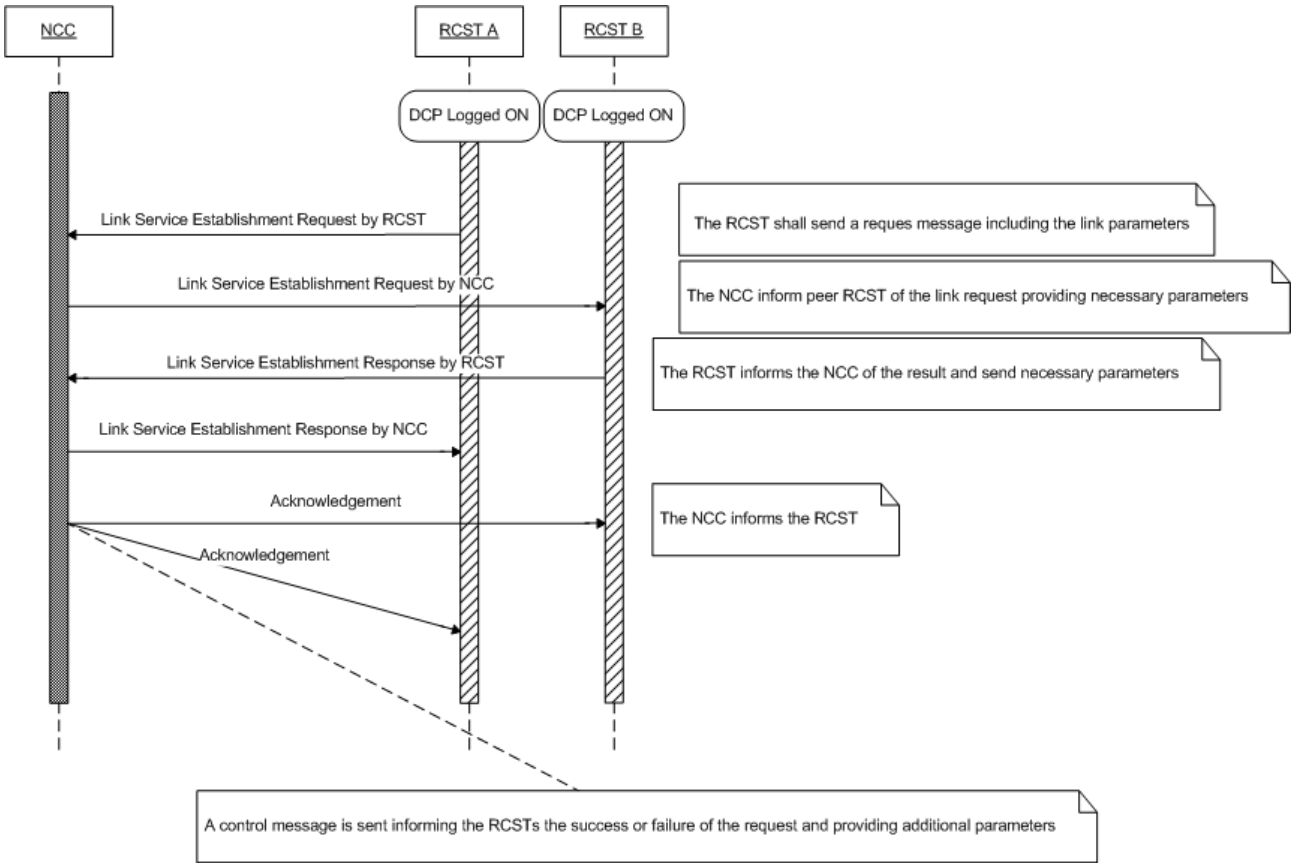


**Figure 13.3: Successful unicast connection**

#### 13.1.2.2 RCST DCP successful multicast connection

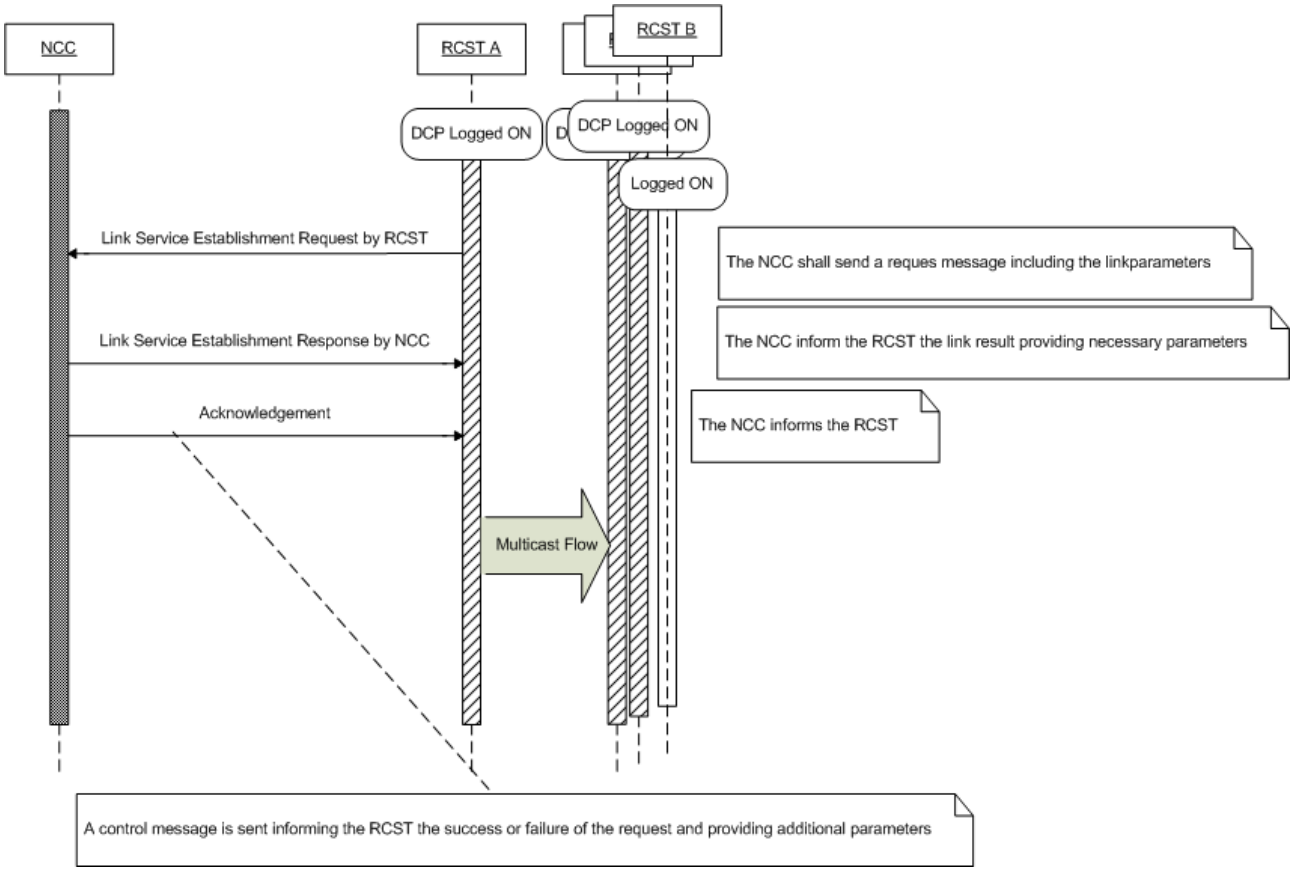


**Figure 13.4: Successful multicast connection**

## 13.2 DCP-enabled RCST state machines

A DCP-enabled RCST should implement at least the DCP Logon and unicast/multicast connection setup procedures. Following clauses describe these state machines. The RCST should also be assigned in its MIB the configurable timeouts and number of retries expressed in the figures.

### 13.2.1 DCP logon

Figure 13.5 shows the state machine for DCP logon. If a DCP LOGON RESPONSE is not received in the wait_for_Response state, the RCST may retry the request for a configurable number of times. The RCST should go to the initial state after a configurable Timeout_T0.

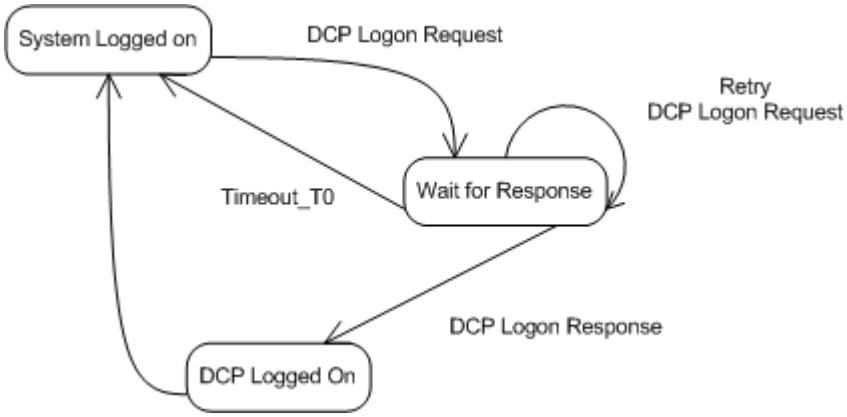


**Figure 13.5: DCP logon state machine**

## 13.2.2 DCP unicast connection

Figures 13.6 and 13.7 show state machinesrun fora connection establishment sequence for the initiating terminal and the peer terminal, respectively. The number of retries and configurable timers are expressed in the figure and should be assigned appropriate values in the RCSTs MIB.

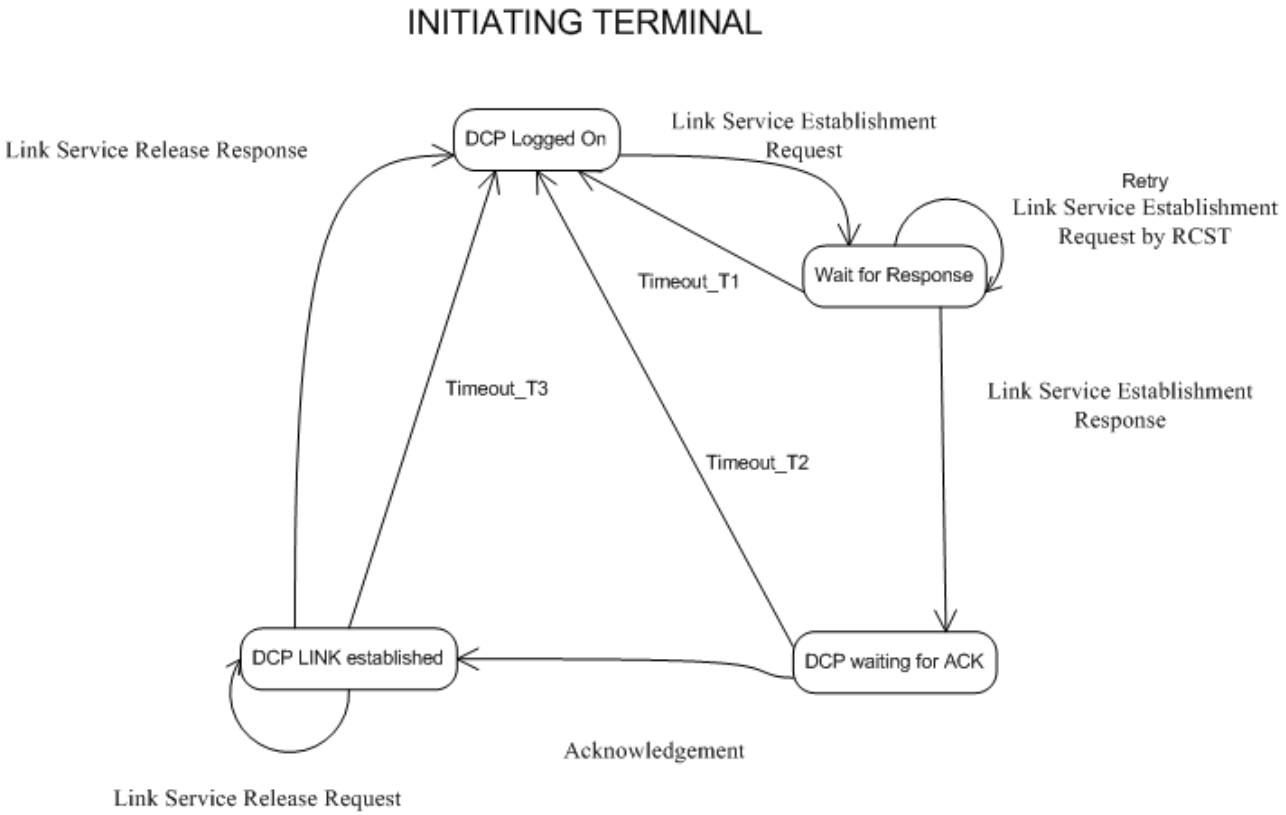


**Figure 13.6: Unicast connection setup in the initiating terminal**

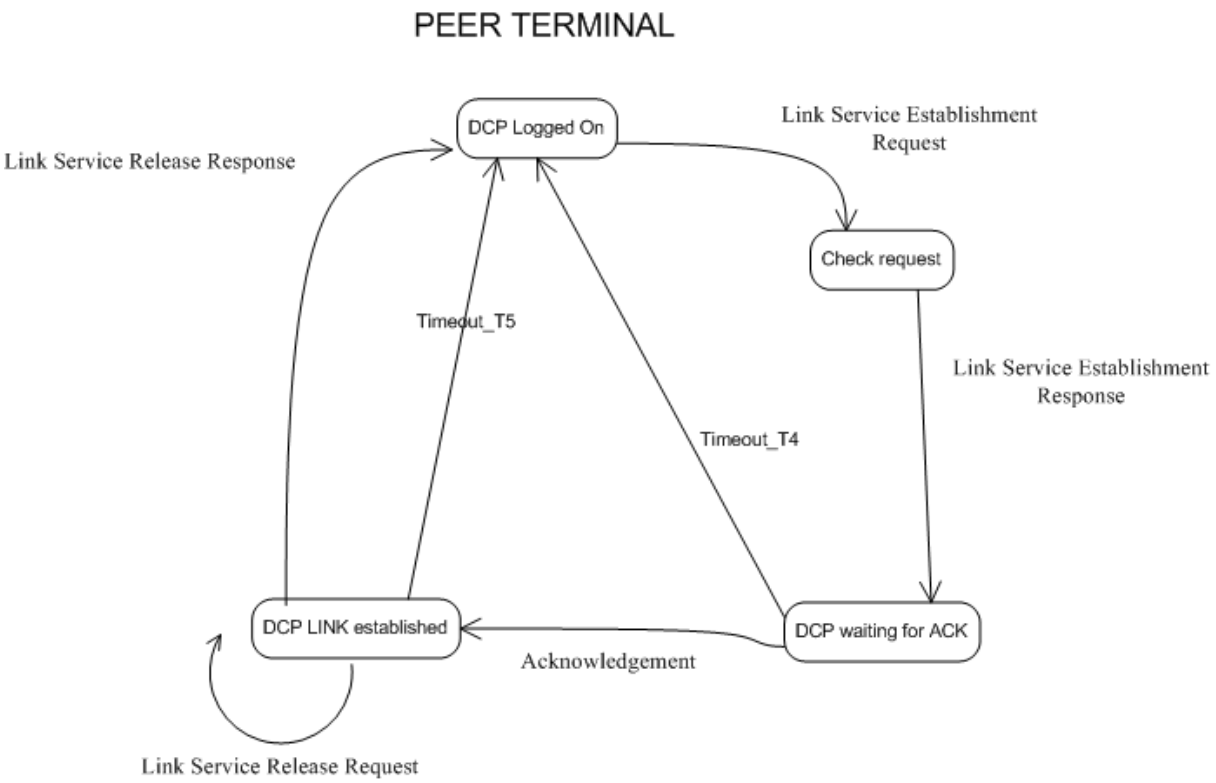


**Figure 13.7: Unicast connection setup in the peer terminal**

## 13.2.3 DCP multicast connection

The multicast connection setup state machine is similar to the unicast connection setup state machine with some different procedures and parameters. This state machine is shown in Figure 13.8.

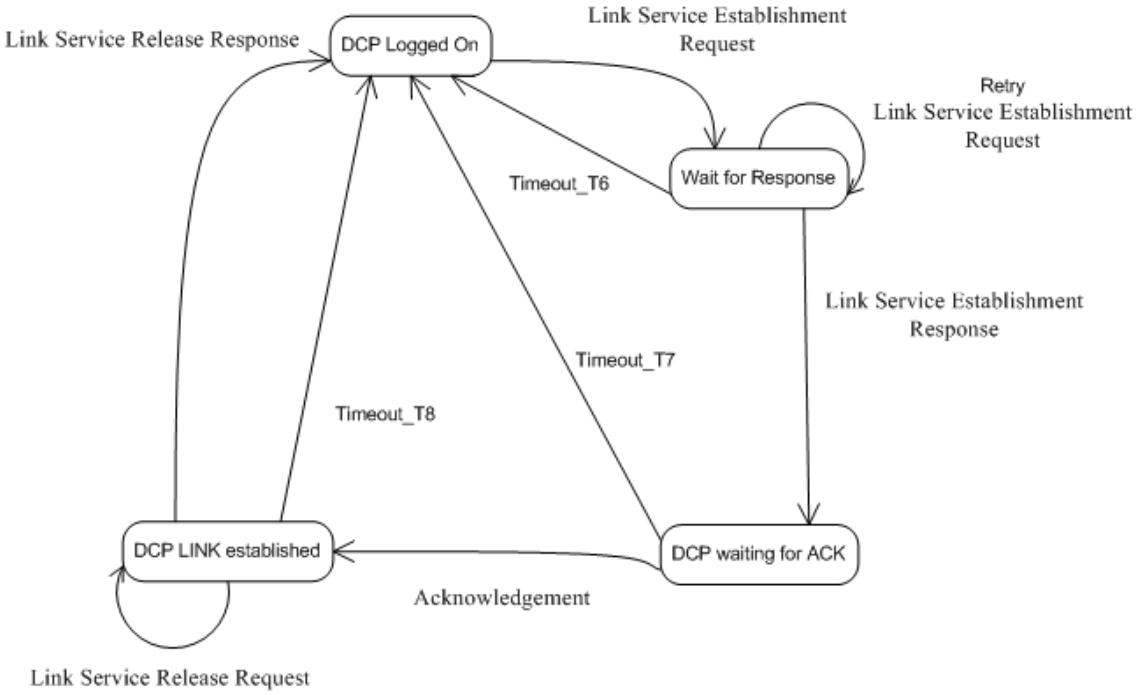


**Figure 13.8: Multicast connection setup in the initiating terminal**

### 13.2.4 DCP routing procedures

The DCP protocol performs unicast/multicast address resolution and routing functions, specifically for meshed systems. If the next hop IP address of an outgoing packet is not found in the AR database, a DCP connection establishment request is triggered by the RCST to find the L2 address of the next hop. In case the system does not support the dynamic routing function (e.g. OSPF), the DCP protocol can assist the RCST with IP routing information.

The NCC does not allow DCP connections across different SVNs or VRF domains.

The RCST may indicate in the request message the next hop IP address (Next hop address field in the Triggering datagram identifier IE) according to its RIB. When this field has been filled by the RCST and the NCC cannot identify the destination RCST from the triggering packet destination address, then the NCC should use the address of the next hop field to obtain the MAC24 address and the FPDU identifiers corresponding to the peer RCST.

### 13.2.5 Other possible DCP functionalities

DCP is also a complement to the functionality of the interfaces already defined in the DVB-RCS2 and DVB-S/S2 standards. Other functions that may be added by the DCP protocol for DVB-RCS2 control plane can be summarized as:

- QoS-driven dynamic allocation of bandwidth resources connectivity channels, following the execution of a Connection Admission Control (CAC) function.
- Dynamic control of the communicating parties in the DVB-RCS2 system, via configuration parameters and policies.
- Dynamic allocation/assignment of logical resources to allocation channels.
- Definition of isolated and independent satellite sub-networks within the global interactive network (i.e. each subnetwork is characterized by its own terminal population, bandwidth resources, addressing space/plan).

# 14 Transparent mesh overlay networking

Table 14.1 presents some terminology specifically used in this clause.

**Table 14.1: Additional terminology for transparent mesh overlay networks**

| Term | Definition |
|---|---|
| Link | In the IP communication context this refers to a *data link* as a sub-IP Connection that can be used for submission of IP traffic destined to a specific range of IP addresses; in the satellite communication context a connection via satellite connecting parts of the ground segment. |
| Behaviour Aggregate | Traffic aggregate that gets unified treatment regarding the transport over a link. |
| Link QoS Class | Traffic classification recognized by the RCST. |
| Link Service | The set of policies used to implement a certain *Link Behaviour* or *Link Behaviour Group* for a specific link. |
| Link Stream | Stream of consecutive PPDUs over a link carrying a consecutive stream of ALPDUs. |
| Link Behaviour | The characteristics of a *Link Behaviour Aggregate* related to the transport over the link. |
| Link Behaviour Group | A set of *Link Behaviours* that have specific common policies related to the transport over the link. |
| Receiver Physical Layer Segment | A part of the physical layer monitored in its entirety by an associated receiver, as seen from the transmitter side. |
| Request Class | Resource request classification recognized by the MF-TDMA resource controller. |
| Link Interface | A sub-interface of the satellite interface. It can be used to reach a subset of the link receivers that can be reached via the satellite interface. |

Dynamic Connectivity Protocol (DCP), which is specified in [i.1], aims to support mesh overlay networking in combination with star topology in the same network.

Five link types can be identified in the applicable networks:

1) Bi-directional link between mesh RCSTs using MF-TDMA
2) Unidirectional link from an RCST to the hub GW using MF-TDMA
3) Unidirectional link from the hub Feeder to an RCST using TDM
4) Unidirectional link from the hub Feeder for transport of multicast with the TDM
5) Unidirectional link from an RCST for transport of mesh multicast using MF-TDMA

The link types 2 and 3 are the existing link types used in forward and return direction in a star DVB-RCS2 system. The link type 1 is a mesh link where a mesh RCST is not only able to send MF-TDMA bursts, but it is also capable of receiving them. Link types 4 and 5 are additions that support multicast.

The main extensions to DVB-RCS2 to provide mesh overlay networking are identified here:

- Each mesh-capable RCST is equipped with a DVB-RCS2 compatible burst receiver, possibly a wideband multi-carrier receiver, to receive MF-TDMA burst transmissions from other mesh RCSTs. This receiver operates concurrently with the DVB-S2 compliant receiver for the reception of the TDM Forward Link from the Feeder.
- The router within each RCST is extended to support IP routes for use within its Mesh Satellite Subnet.
- The RCST supports DCP client part, enabling on-demand mesh link establishment.
- The Network Control Centre is enhanced with a Mesh Controller, which implements the server part of the DCP, and which is responsible for mesh link management and control, as well as mesh routing.

Full mesh networking is created by enabling mesh links between RCSTs on demand. This demand is traffic initiated and it is expressed as a link request to a Mesh Controller containing sufficient information so that the Mesh Controller can identify the correct link destination, IP hosts reachable via this destination and the applicable link service specification. The mesh RCSTs need to be registered or logged on to the Mesh Controller in order to send the link request. How the request will be treated by the Mesh Controller depends on the destination, whether it can be reached over a mesh link, the receiver state, the state of the link service (permanently or temporarily blocked or not) et cetera.

After receipt of a valid link request (no erroneous data, the receiver up and logged on), the Mesh Controller will attempt to establish the link service for the opposite traffic direction as applicable for the link service requested. The establishment of the opposite direction should at least have progressed to either the "link service established" or "link service blocked" state before the initiating traffic arrives at the destination RCST.

## 14.1 Networking principles

A mesh capable RCST first logs on to the NCC as a DVB-RCS2 RCST and then logs on to the Mesh Controller as a DCP client. The RCST provides its IP addresses applicable for mesh routing in the DCP logon request. The DCP server in the Mesh Controller includes this information in a common mesh routing table for a Mesh Satellite Subnet within the Super-frame. If the set of attached IP addresses changes, the RCST will clear all dynamic links and routes, and will logon to the MC again. A new DCP logon will update the MC routing entries related to the RCST.

The permanent hub links and a default route to the hub are automatically established based on the DVB-RCS2 level logon response. This is sufficient to achieve star operation. A mesh capable RCST will only send a DCP logon message if the hub indicates mesh capability. A mesh capable RCST that either connects to a hub that does not support mesh networking or does not get response to its DCP logon message, will map all the satellite traffic to the permanent hub links. In the latter case, the RCST will reattempt DCP logon. The mesh default GWs, the IP address space to be accessed via the hub and the IP address space for mesh networking are set according to the DCP logon response, which occurs as a response to a DCP logon request issued any time after the DVB-RCS2 logon. The mesh default GW may be another mesh RCST and thus a dynamic link may be required to reach the mesh default GW; or it may be at the hub in which case the link to the default GW is permanent.

An IP packet neither identified as mesh traffic nor identified as hub traffic is mapped to the mesh default GW (which can be either the hub or another mesh RCST). An IP packet for the satellite network maps either to a permanent hub link or to a dynamic link. Forwarding of traffic to the hub does not require link establishment control signalling. Temporary rerouting to hub may occur, and this requires the routing part of the link establishment control signalling.

If an incoming packet is identified as mesh traffic it will eventually be mapped to a specific dynamic link. If a mapping to a specific dynamic link is not yet known, DCP is employed to establish the specific route entry and the specific link that will be used for the packet. The latter can be by association to an already established link, and the link establishment is then skipped.

Packet forwarding may be rejected by the Mesh Controller and this packet and similar packets will then be discarded or blocked for a given period (proper ICMP could then be given to the source IP address). After this blocking period a similar packet will again trigger a link establishment attempt. The minimum attempt hold-off period to be used in this case is indicated in the logon response message, as one of the DCP system parameters. The link response message contains one of the link rejected values (since there can be several reasons for link reject) in the reason field.

Traffic may be temporarily rerouted by the Mesh Controller, and the triggering packet and similar packets will then follow the new route for the lifetime of the route. After the expiration of the lifetime, the route will be cleared and a similar packet will again trigger a link establishment attempt, allowing a renewed routing decision.

The Mesh Controller will provide explicit routing information to avoid that e.g. NMS traffic goes to the mesh default GW, when the default GW is a mesh RCST. This is achieved by informing the mesh RCST of the destination address space that applies to the hub link, so that the mesh RCST sets up permanent mapping to the permanent hub link for this address space. The hub can be the mesh default GW, or any mesh RCST can be used as mesh default GW. The applicable mesh default GW is identified by using the DVB IP address of the mesh RCST or the hub, as applicable. The RCST will set up a dynamic link to the mesh default GW when required. It is also possible to identify a secondary mesh default GW, which will be used in case the primary default GW cannot be reached.

Traffic to a mesh default GW with an IP address outside the mesh address space and outside the local LAN address space is mapped to the permanent hub link, independent of the hub link address space. This link is the default route when no specific link can be found. A mesh default GW specification within the mesh destination address space (possibly, the user traffic interface address of the RCST that acts as mesh default GW) implies that a dynamic link is required to reach the mesh default GW. Further, the mesh default GW can be set to an IP address of the local LAN address space, used specifically at the RCST that acts as mesh default GW. This default GW IP address should be the address of the LAN side default GW to get into the WAN.

The RCST should not default to route packets received from the MF-TDMA onto the MF-TDMA again.

The assumed next-hop address is indicated, if known, in the link establishment request, if this address is already known to the RCST. If not, it will be set to all-zeros. The DCP server should not accept link establishment to the next-hop address if the next-hop address and destination address do not match the current network topology as known by the Mesh Controller. The RCST should then be instructed to clear the erroneous route entry.

The RCST will clear any existing route entries which conflicts with the route entries conveyed to the RCST through the most recent call establishment signalling. (Note that overlapping routes may exist).

## 14.2 Mesh multicast

The mesh multicast addresses a RCST is allowed to forward, if any, are given in mesh Logon Response, as a part of DCP address space. RCST adds all received mesh multicast address to its multicast routing table when parsing Logon Response message.

If a change in mesh multicast configuration occurs, RCST will be notified by DCP server using Link Service Establishment Request message with status/reason field indicating change of mesh multicast configuration (see table 9 for coding of the status). Mesh multicast addresses configured for the RCST will be given in Route Entries for link IE of the Link Service Establishment request. Content of other IEs of the message will not be applicable. Upon reception of this message, RCST will update its DCP address space and its multicast routing table by adding new multicast address and removing addresses no longer applicable.

DCP server will repeat sending this LSE request message with mesh multicast change status until it receives a response or a number of outstanding responses has reached its limit. Link ID for this message will always be 0, while the service id will be the message identifier that RCST will have to use when replying to the message.

Traffic to a mesh multicast address received at LAN side will initiate mesh link establishment, if an applicable route cannot be found, and the link will be granted if the QoS class the link is requested for is configured for the RCST. If not, link will be denied. TXID is assigned and may be the same as the TXID assigned for unicast.

Update of mesh multicast configuration is not sent to RCST if there only has been change in the QoS class configured for the RCST. In that case, DCP server will release any existing mesh multicast link that is up and cannot longer be supported, and it will not allow establishment of a new link on not supported QoS link.

On the receiver side, there is no mesh link establishment related to reception of mesh multicast traffic. Mesh multicast addresses and TXID are broadcasted in the MMT table, encoding the TXID in the elementary_PID field. On reception of IGMP join from LAN side, the RCST may open the receiver for the relevant TXIDs as required, if any is found in the MMT table for the subscribed multicast.

# 14.3 RCST MF-TDMA transmitter

## 14.3.1 RCST protocol architecture

Figure 14.1 illustrates with an example the principal protocol structure of a mesh RCST.

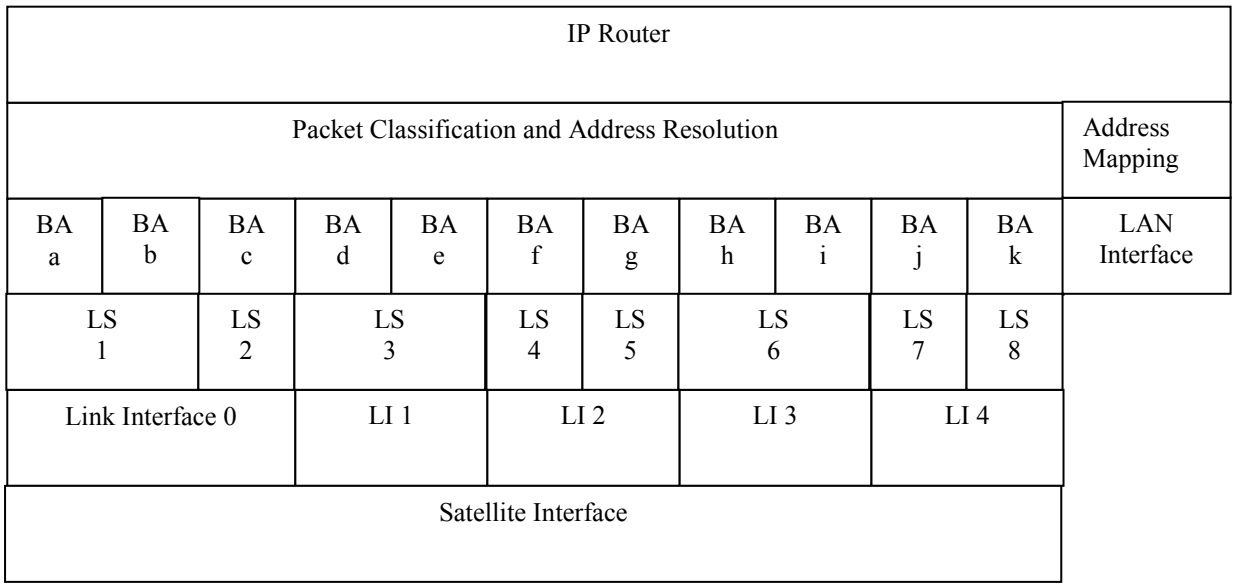


**Figure 14.1: Structure of the mesh RCST interfaces, showing concurrent link streams at each Link Interface**

The satellite interface is in the example in Figure 1 divided into 5 link interfaces (LI). Each LI supports here one or two concurrent link streams. Within a LI, there is one BA for each QoS. A BA maps to one LS of the associated LI. Several BAs using the LI may share a link stream (LS) as a single SA, or a solitary BA may have a dedicated LS. This is a policy choice at the transmitter side. There may be several BAs in use towards a mesh RCST at the same transmitter.

LI0 connects to the hub. The associated BAs and LSes are for this LI automatically set up at DVB-RCS2 logon, independent of the MC and without involving DCP. The mesh RCST uses DVB-S2 TDM for LI0 reception and DVB-RCS2 MF-TDMA for transmission.

LI1-LI4 each represents a satellite link to one or several mesh RCSTs in the mesh satellite subnet. This segregation is non-overlapping. These links are based on DVB-RCS2 MF-TDMA in both directions. Associated BAs and LSs are set up and released controlled by DCP signalling between the RCST and the MC.

The LL Service (LLS) serves a single Link Behaviour (LB) or an LB group. The LLS is subject to certain policies of which some are inherited from the HL association, some are permanent, some are configured in advance, some are signalled through DCP and some are enforced by the resource controller.

A system may use a separate Receiver Physical Layer Segment (RPLS) per satellite link destination or it may be based on sharing RPLS between several link destinations. The transmitter may in the latter case use shared BAs and shared LSes, merging the traffic aggregates for two or more destinations having the same RPLS into one LI. It can be assumed that any timeslot applicable for the RPLS can be used to reach the destinations monitoring the RPLS given that sufficient power and waveform can be used in the timeslot. An Assignment ID value points to only one RPLS. The use of timeslots with Assignment IDs pointing to different RPLS cannot be swapped between LIs.

Traffic is sent over the mesh LIs as DVB-RCS2 PPDUs, kept apart between transmitters by use of different FPDU Transmitter identifiers (TXID) and between LSes of the same transmitter by different PPDU Fragment ID values, and kept apart between destination RCSTs by use of individual receiver ALPDU MAC24 addresses. The LS is determined at the transmitter by associating a packet to a LI, and to a BA of that LI. The packets of a BA are all sent in the same LS. The PPDUs of the LS are placed into FPDUs that are transmitted in timeslots that are known to be monitored by the destination RCST, i.e. part of the RPLS monitored by that RCST. SDU reception is achieved by monitoring applicable timeslots, reassembling ALPDUs from each LS that may be applicable for the receiver, dropping the ALPDUs not aimed for the receiver and making a forwarding decision for each SDU that is aimed for the RCST.

The Request Class (RC) is used towards the resource controller in the NCC to identify the LLS associated with each capacity request, and the required RPLS. The resource controller indicates through the Assignment ID each timeslot assigned to the mesh RCST transmitter. The Assignment ID values are used to segregate the timeslot resources for each LI. Segregation of LQC may be done using the Assignment ID. Alternatively, the RCST may be allowed to map each timeslot to a pool for each LI and utilize for each LI these timeslots according to the applicable QoS policies.

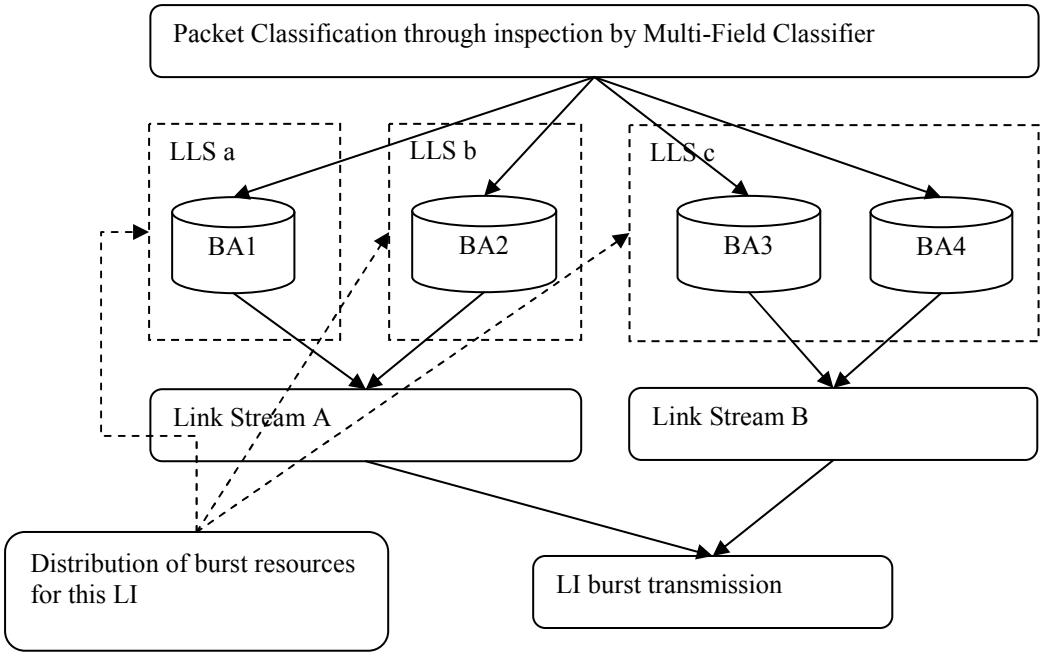


**Figure 14.2: Internal Link Interface packet classification and transmission scheduling**

Each LI supports a number of queues, each used for a specific BA, as illustrated in Figure 14.2. Each BA maps to one LLS. The LLS may serve an LB group by applying different policies for the BAs it is serving.

Figure 14.2 illustrates that the distribution of the timeslots controls the shaping and scheduling of submission from each LS and between the LSes. PPDUs from different LSes are interleaved onto the LI. SDUs from different queues are interleaved in a shared LS. The BA queues are FIFO queues unless packets have to be dropped from the queue due to over-filling.

The FPDUs of the different link LIs are interleaved into the TX Stream. Each LI maps to only one RPLS. The RCST assumes that each Assignment ID value consistently identifies a single RPLS. *All* the timeslots of an RPLS are assumed monitored by the recipients indicated to monitor this RPLS. An RPLS may be as narrow as one single RCST or as wide as all the carriers in use by the super-frame, or a subset of the carriers in the super-frame.

## 14.3.2 Routing

The satellite sub-interface address resolution table in the terminal is populated with static elements by the mesh controller through the DCP logon response, separating between default address resolution and dynamic address resolution. The static entries are given by the DCP address space and hub links address space received in DCP Logon Response. The mesh controller populates the address resolution table further with dynamic entries on demand. The dynamic entries may map to static and dynamic links.

Table 14.2 shows an example of a satellite interface side routing table at RCST.

**Table 14.2: Example of RCST routing table for a transparent mesh overlay network**

| # | Base address | Mask | Next hop | Virtual link | Metric | Comment |
|---|---|---|---|---|---|---|
| 1 | 10.12.0.0 | 255.255.0.0 | - | 0 | 0 | Star address space; Inserted by DCP logon response |
| 2 | 10.13.0.0 | 255.255.0.0 | - | 0 | 1 | Star address space; Inserted by DCP logon response |
| 3 | 10.10.0.0 | 255.255.0.0 | - | - | 0 | DCP address space; Inserted by DCP logon response |
| 4 | 10.11.0.0 | 255.255.0.0 | - | - | 0 | DCP address space; Inserted by DCP logon response |
| 5 | 10.10.11.0 | 255.255.255.248 | 10.10.11.01 | 4 | 1 | Tied to an active mesh link; cleared when services are cleared |
| 6 | 10.10.85.0 | 255.255.255.248 | 10.10.11.01 | 4 | 2 | Tied to an active mesh link; cleared when services are cleared |
| 7 | 10.10.12.0 | 255.255.255.248 | destination temporarily unreachable | - | - | Hold-off to block temporarily in order to limit signalling |
| 8 | 10.10.13.0 | 255.255.255.248 | temporary route to the hub | 0 | 3 | Hold-off to forward to the hub temporarily for connectivity |
| 9 | 0.0.0.0 | 0.0.0.0 | 10.10.10.100 | 3 | 4 | Primary default route used for destinations outside Mesh and Hub address space |
| 10 | 0.0.0.0 | 0.0.0.0 | 10.10.10.110 | | | Secondary default route; for redundancy |
| 11 | 0.0.0.0 | 0.0.0.0 | 10.10.10.120 | | | Alternate secondary default route; for redundancy |

Entries 3-4 define non-contiguous ranges in the address space where DCP will be used to resolve destination to next hop. They are inserted by the DCP logon response. Entries 5-8 are inserted by dynamic DCP link control as a result of DCP link establishment. Entries 9-11 define alternate default routes to default GWs. They are inserted by the DCP logon response. They are resolved to link through the associated next hop IP address (here 10.10.10.100 for the primary default GW).

Entries 5 and 6 are routes for active mesh link services. The entries are removed when the respective link services are released. The routes use the same virtual link as they go via the same next hop router. The virtual link identifies the RPLS monitored by the receiver of the next hop router.

Entry 5 describes the destinations directly attached to the link receiver with no router in-between. Entry 6 describes a subnet that is behind another router. It is feasible that the entry 4 subnet could also be reached over another link with lower metric.

Entry 7 is established due to a link service rejection and the cause is failure to connect to the destination RCST. The status is then that the next-hop is neither reachable over the TDMA nor via the default route. Other routing possibilities to the destination are not known. The entry is cleared at blocking timeout. This type of entry can limit useless DCP signalling and useless transmission to the hub.

Entry 8 is established due to a mesh link rejection and the cause in the link establishment response. The planned next-hop is reachable but not by direct mesh link. It is believed to be reachable via the hub. The entry is cleared at blocking timeout. The reason for this type of reroute may e.g. be current link conditions and service policies (would require service class based address resolution).

The metric will have to be larger than 2 to allow several metric values to be used for differentiated routing over multiple mesh links.

Entry 9 is the default route to the default GW. This route is never cleared. It may map to a virtual link that uses DCP link service establishment. It may map to the hub and then the necessary virtual link and corresponding links are automatically established at RCST logon.

Entries 10 and 11 are routes to secondary default GW.

### 14.3.3 Link and Link Service establishment and release

#### 14.3.3.1 Establishment

A LL service is established for a specific link. In cases where there is only one LL service for a link, the LL service establishment is integrated with the link establishment. DCP is employed to establish the specific route entry and the specific link that should be used for the packet mapping to this LL service. The need to differ link from LL service establishment comes from the fact that there might be several LL services using the same link (transmitting packets of different QoS classes between two RCSTs). The first time we establish a link service between two RCSTs, also the link is established. In the response from the Mesh Controller a global reference is provided for the link, and the Request Class identifier to use when asking for capacity for this specific link service. Also, the MC provides routing information applicable for the link, and the TXID to be used by the transmitter. If needing to establish another link service between the two RCSTs, neither the global reference for the link nor the link routing information is required, since we already have it, but request class for requesting capacity and TXID to use (this may be the same TXID that is already in use for another link service). So, the process of link service establishment will be the same in both cases. The difference is that the Mesh Controller in the second case will not assign a new global link reference; it will duplicate the link establishment response from the first case apart from giving a new LL service ID, a new request class and maybe a new TXID. An initial LL service ID is chosen by the RCST in the link establishment request. The Mesh Controller will in its response to the RCST assign a LL service ID for that specific service. This ID is unique within the link it is established for (within the link identified by the global link reference). It is therefore a unique reference of a LL service within a link as seen from the Mesh Controller and the participating RCSTs. The value of the LL Service ID used in the Link Establishment Request from the RCST indicates the LQC of the requested LL service. In other signals, the LL Service ID indicates the unique identifier for the link service.

The remote RCST will receive link establishment request by the Mesh Controller with all necessary link and link service information (as already decided by the link initiator and the Mesh Controller). It needs either to accept or reject the link establishment.

#### 14.3.3.2 Release

Complementary to link and LL service establishment, there is link and LL service release. A link will exist as long as there is a LL services that belongs to it. When the last LL service is released, the link itself is released. The RCST initiating link establishment will be the main responsible for releasing LL service and links. The remote RCST will also release an inactive link after a long timeout in order to avoid a hang situation. The mesh Controller also has the opportunity to release LL services and links.

## 14.4 RCST MF-TDMA receiver

A mesh-capable RCST needs to be equipped with a DVB-RCS2 compatible burst receiver, and possibly a wideband multi-carrier burst receiver. The receiver needs to be selective of the set of TXIDs that applies for it in its current state. These may be signalled by the means of DCP, signalled through the MMT, or possibly locally synthesized from TBTP2 parsing.

Unicast traffic is specifically addressed to the MAC24 address assigned at DVB-RCS2 logon.

Nominally, a burst is associated to a specific transmitter by the explicit TXID in the FPDU, and the contained ALPDU is associated to a specific receiver through the explicit MAC24, in combination associating the ALPDU to a specific link. The burst receiver may however be built to accept packets without MAC24 arriving in timeslots known in advance to be used for a specific link, even without explicit TXID tagging. Such selectivity is feasible by using receiver timeslot selectivity throughout the system. This can be exploited for header compression. The link specific timeslots should be sufficiently identified by the transmitter and Assignment ID.

The burst receiver is tuned to the RPLS. The RPLS can be built in different ways, and may be synthesized locally in real time from the TBTP2 based on the dynamic information given by DCP. By this design the burst receiver cannot be expected to monitor other slots than those explicitly assigned, but the burst receiver can neither be expected to not monitor other timeslots on the indicated mesh carriers, if such suppression is not specifically known supported by the implementation.

## 14.5 Adaptive Coding and Modulation, and adaptive timeslot sizing

DVB-RCS2 offers the opportunity to do per-burst ACM and adaptive timeslot sizing. A possible strategy for exploiting this may be to:

- Parse the TBTP2 also at the receiver side to determine the timeslot structure of each individual frame, and the modulation and coding that will be used in each individual timeslot.
- Prevent the mesh transmitter from using transmission types on a specific link that cannot be expected to close the link to the specific mesh receiver.

# 15 Dynamic connectivity protocol guidelines for transparent mesh overlay networks

## 15.1 Mesh carrier frequencies

The NCC sends in TIM-B a Mesh System descriptor with a list of frames that may be used for mesh traffic, for each Super-frame used for mesh. The descriptor also indicates the transponder frequency offset for these frames.

## 15.2 Mounting DCP

The DCP messages are transported as UDP packets using a specific DCP UDP port and specific IP addresses, and the NCC needs to indicate this to the mesh RCST.

The Mesh Controller IP address, the DCP multicast address and the UDP port number used for exchanging DCP messages is given in Logon Response TIM-U. An Extension Protocol Descriptor is used to indicate these connection details for an extension protocol. Extension protocols may be used to supplement the lower layer signalling system via IPv4 M&C. The DCP protocol is mounted with the reception of this descriptor.

## 15.3 RCST mesh capability signalling

The DVB-RCS2 logon request message of a mesh capable RCST may inform the NCC of the transparent-mesh capability of the RCST.

## 15.4 DCP message transport

The DCP signalling is UDP/IPv4 based, apart from the initial mesh information exchanged from using DVB-RCS2 L2S signals, as explained in the earlier clause. Table 15.1 gives an overview over the DCP messages. The DCP messages are organized into two main groups, management messages and link control messages.

## 15.5 Summary of DCP messages

**Table 15.1: DCP messages**

| Management Messages | Direction | Description |
|---|---|---|
| DCP Logon Request | RCST → Mesh Controller | Used to log on a mesh RCST to a mesh network; provides RCST terminal and router information to be exploited by Mesh Controller for later link establishments |
| DCP Logon Response | Mesh Controller → RCST | Response from the Mesh Controller; gives the RCST configuration, DCP system information and the hub space and mesh space route information |
| DCP Agent Management Request | Mesh Controller → RCST | Used by Mesh Controller to get the DCP client to perform specific tasks, such as:<br>*clear all dynamic link and routes,<br>*clear all session data and logon,<br>*clear all session data and go to star-only state<br>*leave star-only state and logon to Mesh Controller |
| DCP Agent Management Response | RCST → Mesh Controller | Response by RCST. |
| **Link Control Messages** | **Direction** | **Description** |
| DCP Link Service Establishment Request | RCST → Mesh Controller<br>Mesh Controller → RCST | Either by RCST or Mesh Controller, used to establish link service<br>Also used for mesh multicast configuration update when sent from Mesh controller |
| DCP Link Service Establishment Response | RCST → Mesh Controller<br>Mesh Controller → RCST | Either by RCST or Mesh Controller, response to a link service establishment |
| DCP Link Service Release Request | RCST → Mesh Controller<br>Mesh Controller → RCST | Either by RCST or Mesh Controller, Used to release link service |
| DCP Link Service Release Response | RCST → Mesh Controller<br>Mesh Controller → RCST | Either by RCST or Mesh Controller, response to a link service release |
| DCP Link Status Enquiry | Mesh Controller → RCST | Depending on the System options enabled, could be link quality polling. |
| DCP Link Status Response | RCST → Mesh Controller | RCST response to the mesh status enquiry |
| DCP Link Service Control Acknowledgment | RCST → Mesh Controller<br>Mesh Controller → RCST | Used where just a handshake is required |

## 15.6 DCP message sequence diagrams

### 15.6.1 DCP logon

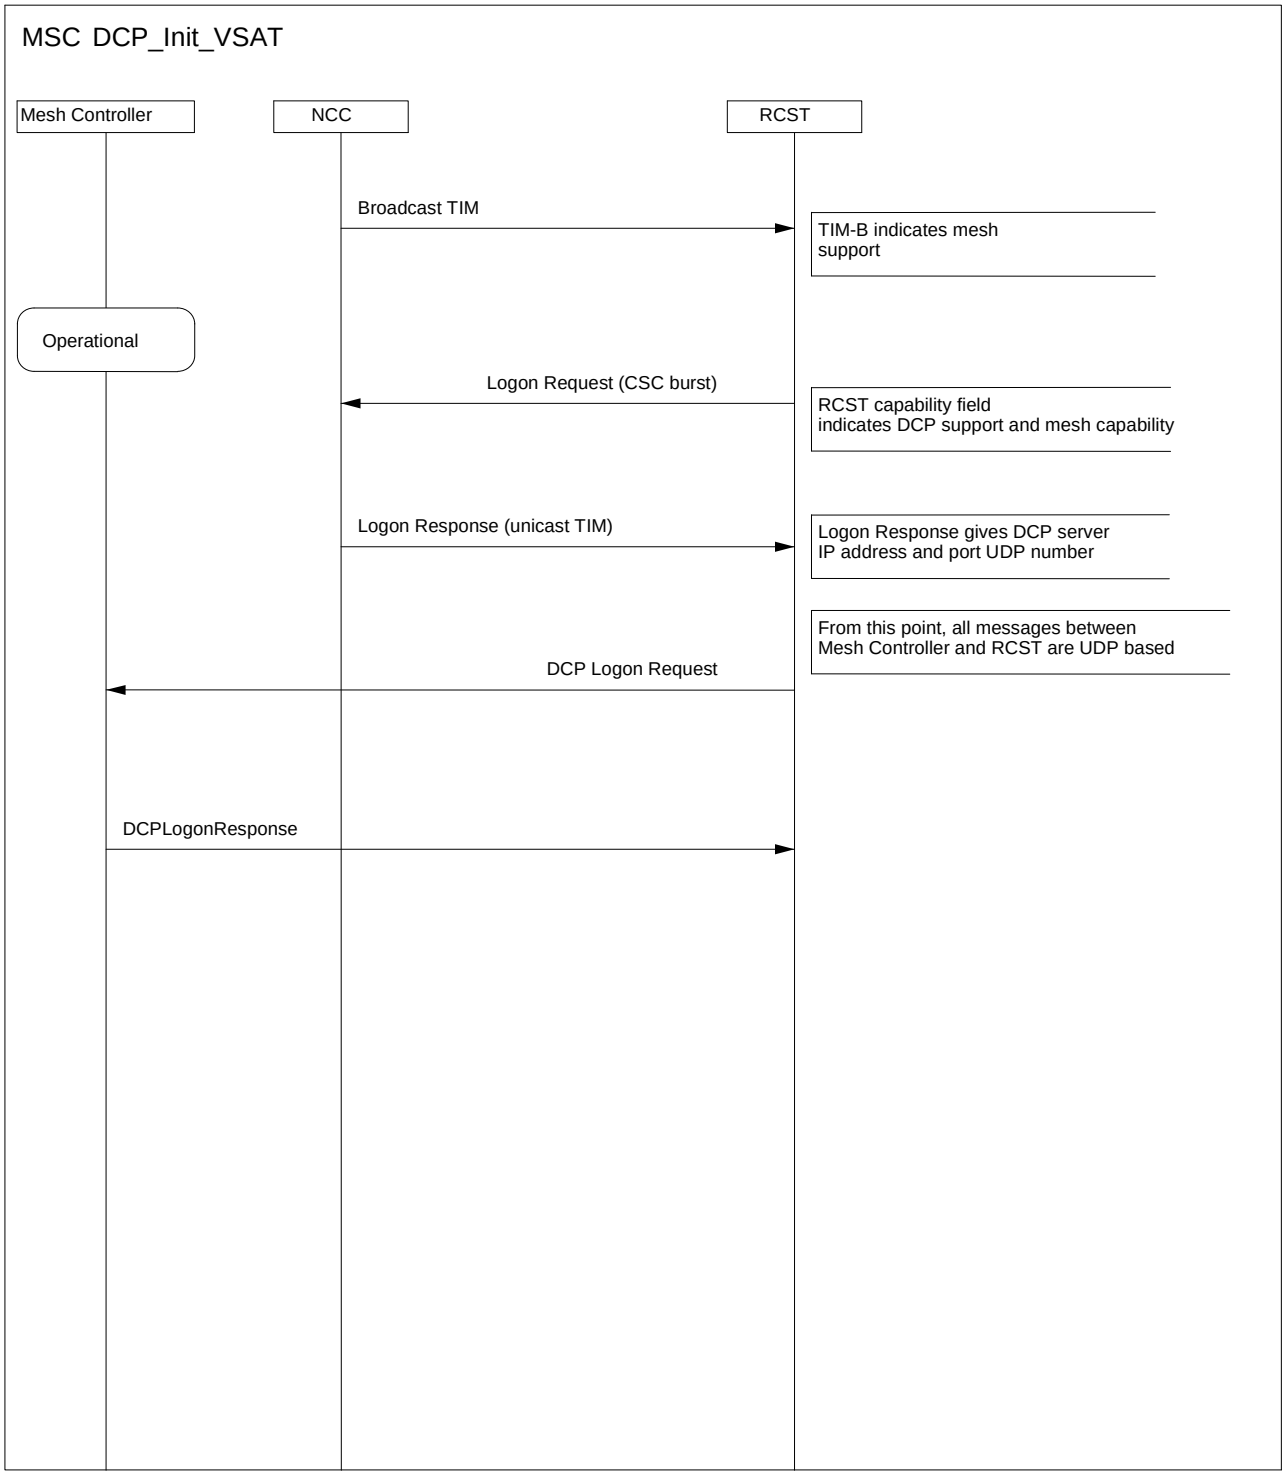


**Figure 15.1: Message exchange during DVB-RCS2 logon and DCP log on - Successful DCP logon**

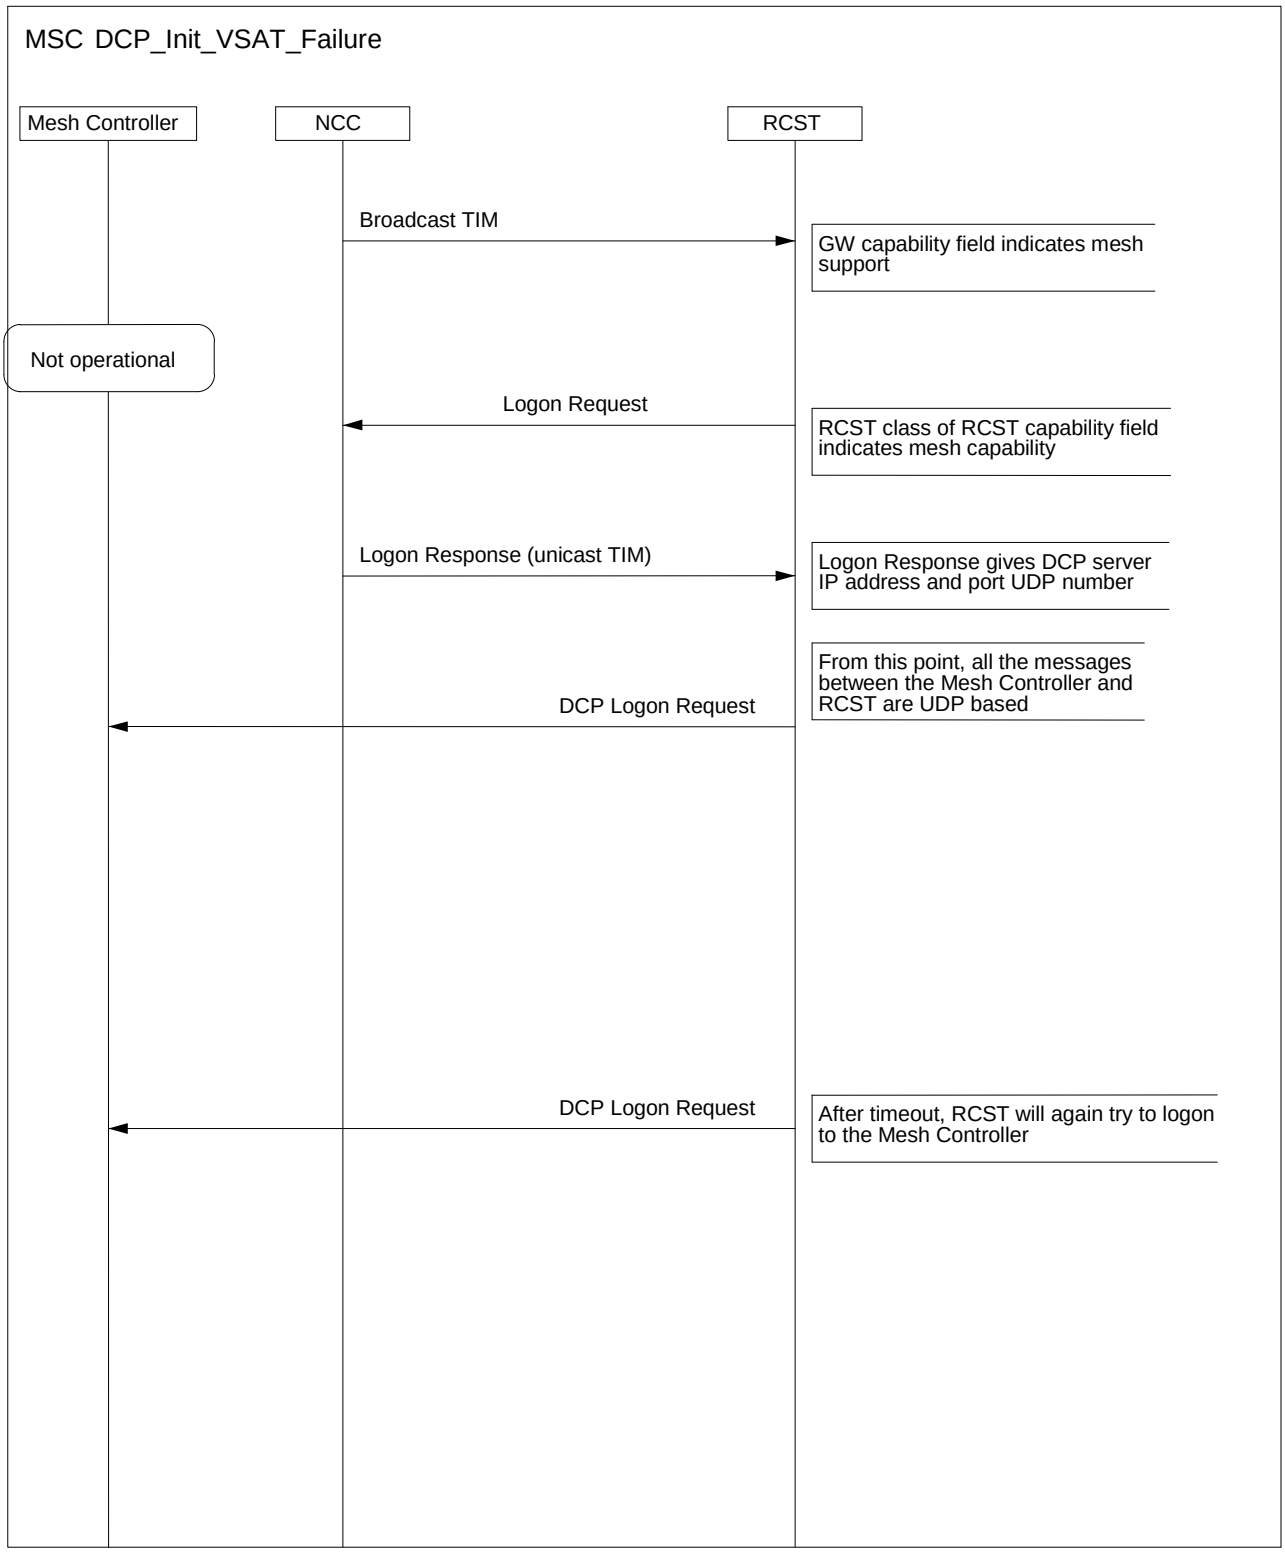


**Figure 15.2: Message exchange during DVB-RCS2 and DCP logon - Unsuccessful DCP logon**

Mesh RCST will continue to send DCP logon requests until it succeeds to log on. The frequency of logon requests should be limited (RCST should keep trying because the Mesh Controller may be expected to come up; if the Mesh Controller has not been intended to come up, the NCC would not have distributed the basic mesh information in the TIM-B and TIM-U).

If a mesh RCST does not receive DCP logon response from the Mesh Controller, it will map all the satellite traffic to the permanent hub links and thus operate in the star mode.

### 15.6.2 Link Service Establishment

Figures 15.3 and 15.4 show examples of link service establishment with two different outcomes.

Re-negotiation of service profile is not supported. If RCST B responds to the Mesh Control Establishment request with the lower service profile then the one it received from the Mesh Controller, the link service will be torn down.

Figure 15.3 also illustrates the capacity assignment kick-start initiated from the Mesh controller that is needed to reduce the initial packet round-trip time. Without capacity requests from the Mesh controller which are sent on behalf of the RCSTs involved, the length of the initial packet round-trip time would inevitably caused retransmissions.

Figure 15.4 shows an example of link service establishment failure. In this case the link service establishment was rejected because the RCST B was not logged on. There are many other reasons for link service establishment rejection. It can be due to an unauthorized request from RCST A, or a request containing erroneous data, it can be due to reject from RCST B because of lack of resources or some other.

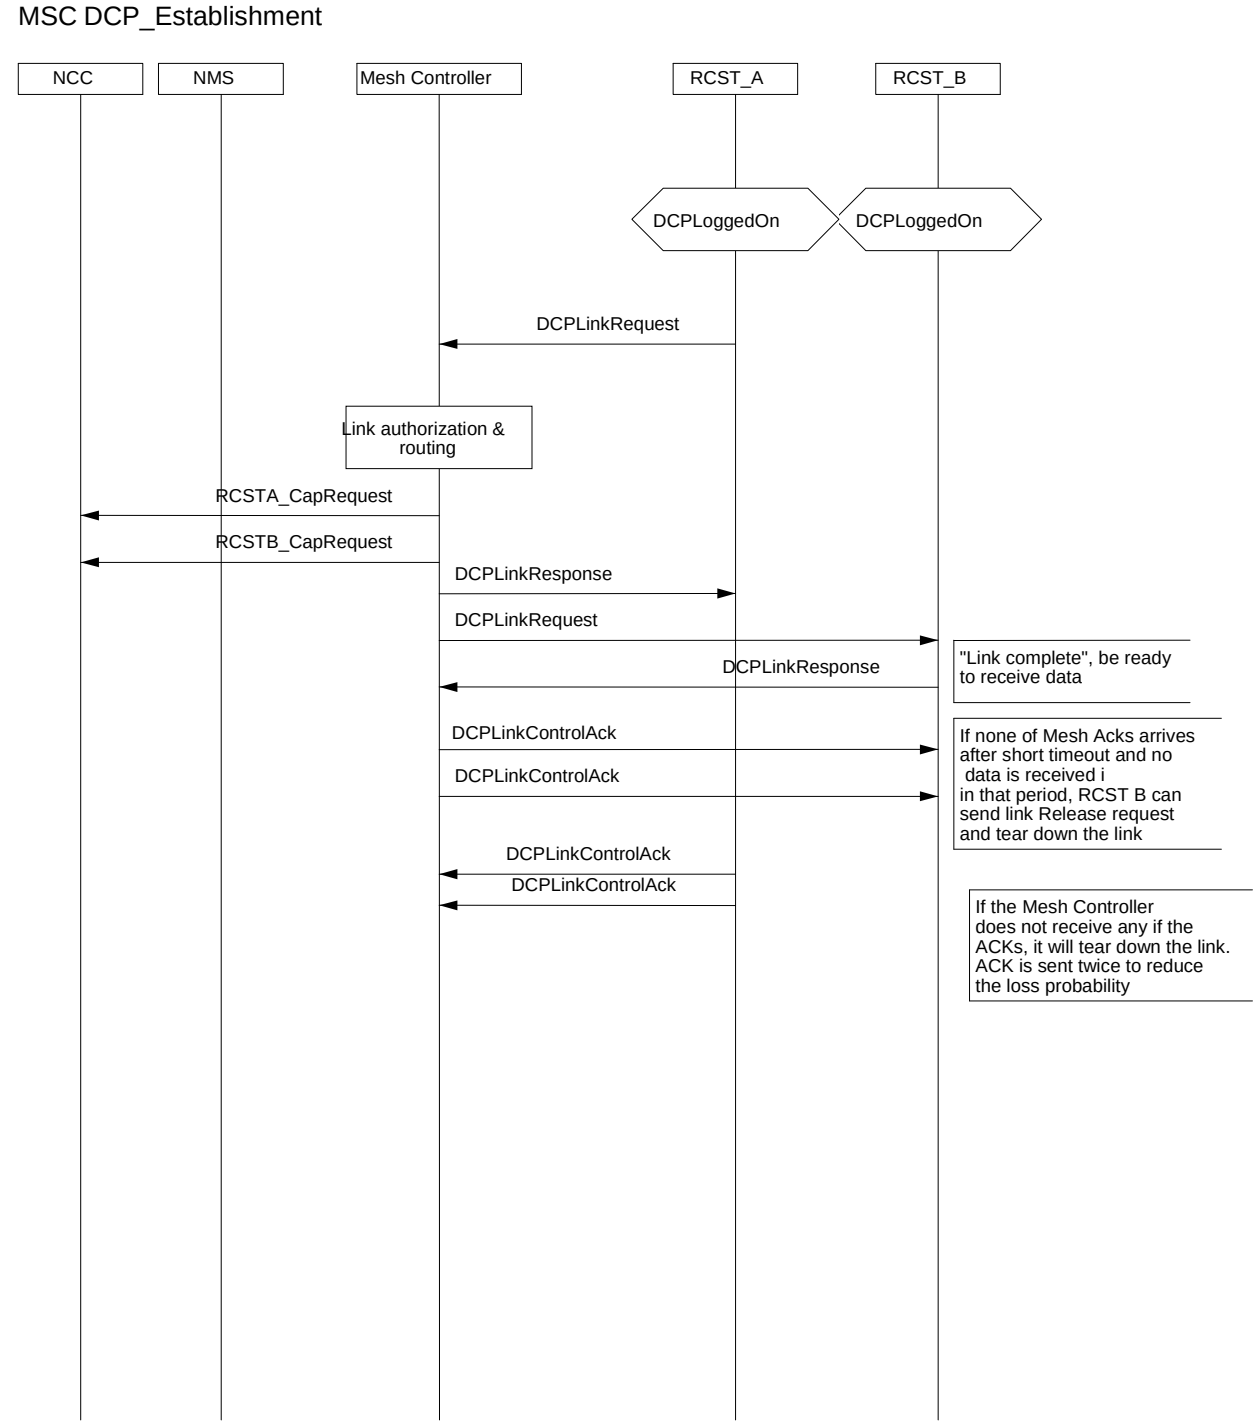


**Figure 15.3: Example of DCP message exchange during link establishment**

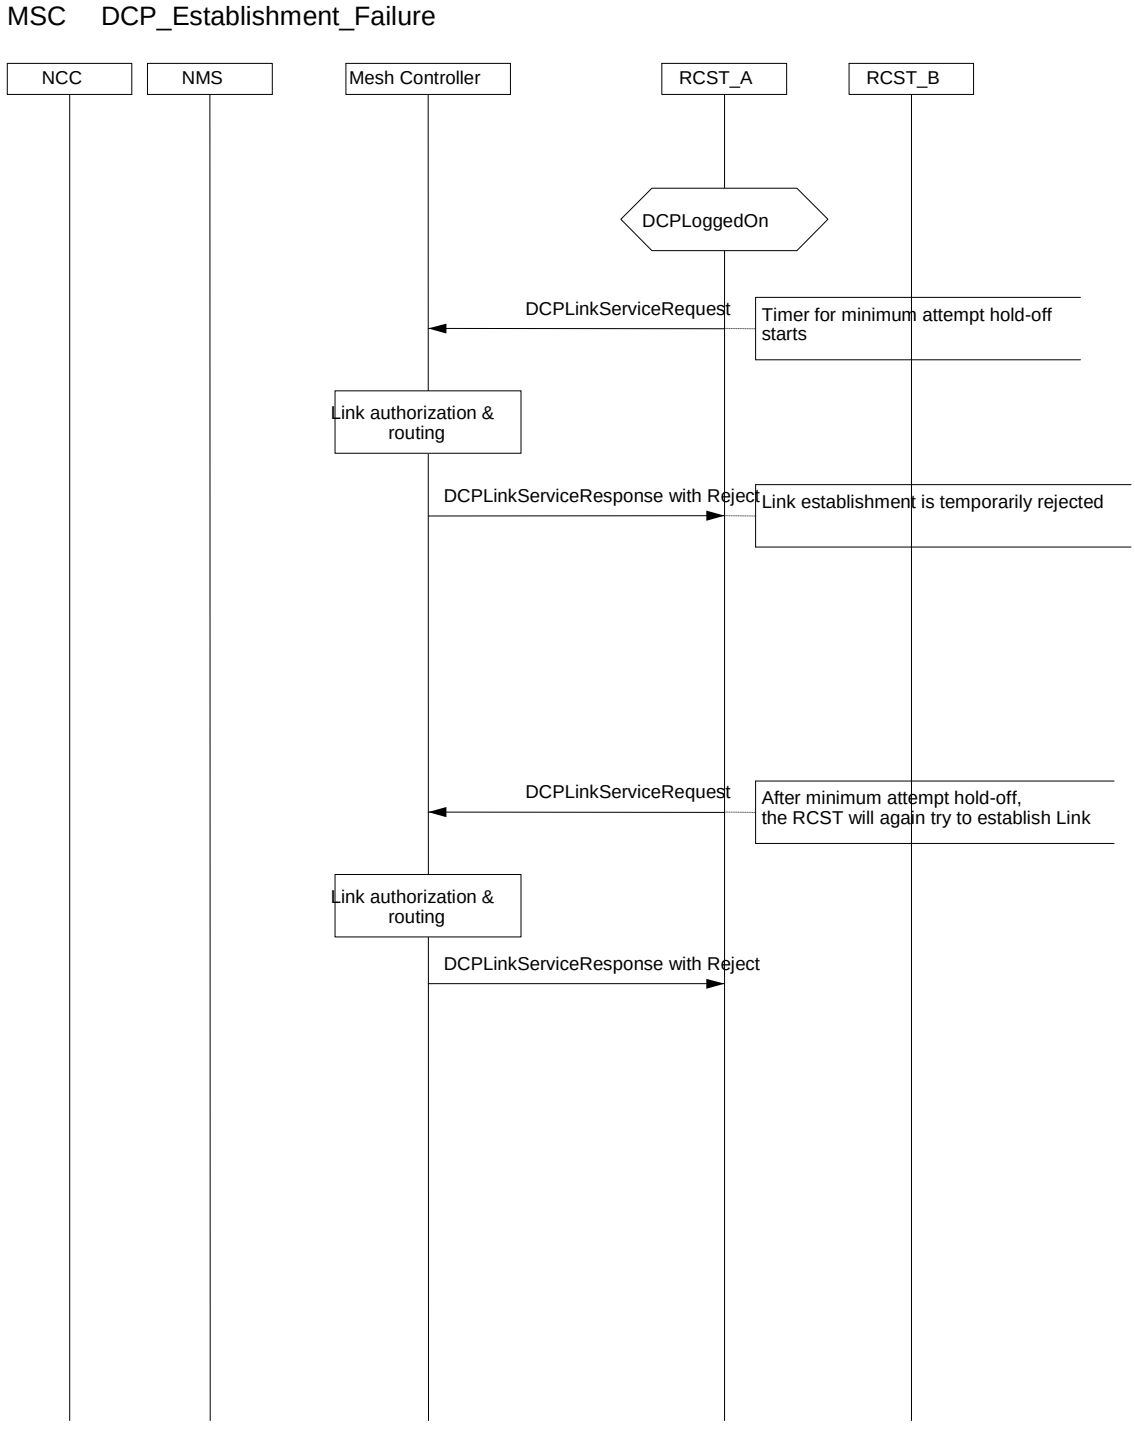


**Figure 15.4: Example of DCP Link Establishment Failure**

## 15.6.3 Link Supervision

A RCST (A) that initiates a link service should also drive the supervision of the connectivity and quality of the corresponding link. It should do this by sending a keep-alive message to the peer RCST nominally 2 times per idle/fixed timeout interval, independent of the level of traffic. The peer RCST should immediately respond with a corresponding supervision message and should restart its peer link supervision timer. If RCST A does not get a response to its keep-alive signal, it should initiate link service release towards the MC. If RCST B does not receive a keep-alive signal before supervision timer expiry it should initiate link service release towards the MC. These internal keep-alive signals should not contribute to keep the connection alive, only externally initiated traffic should do that.

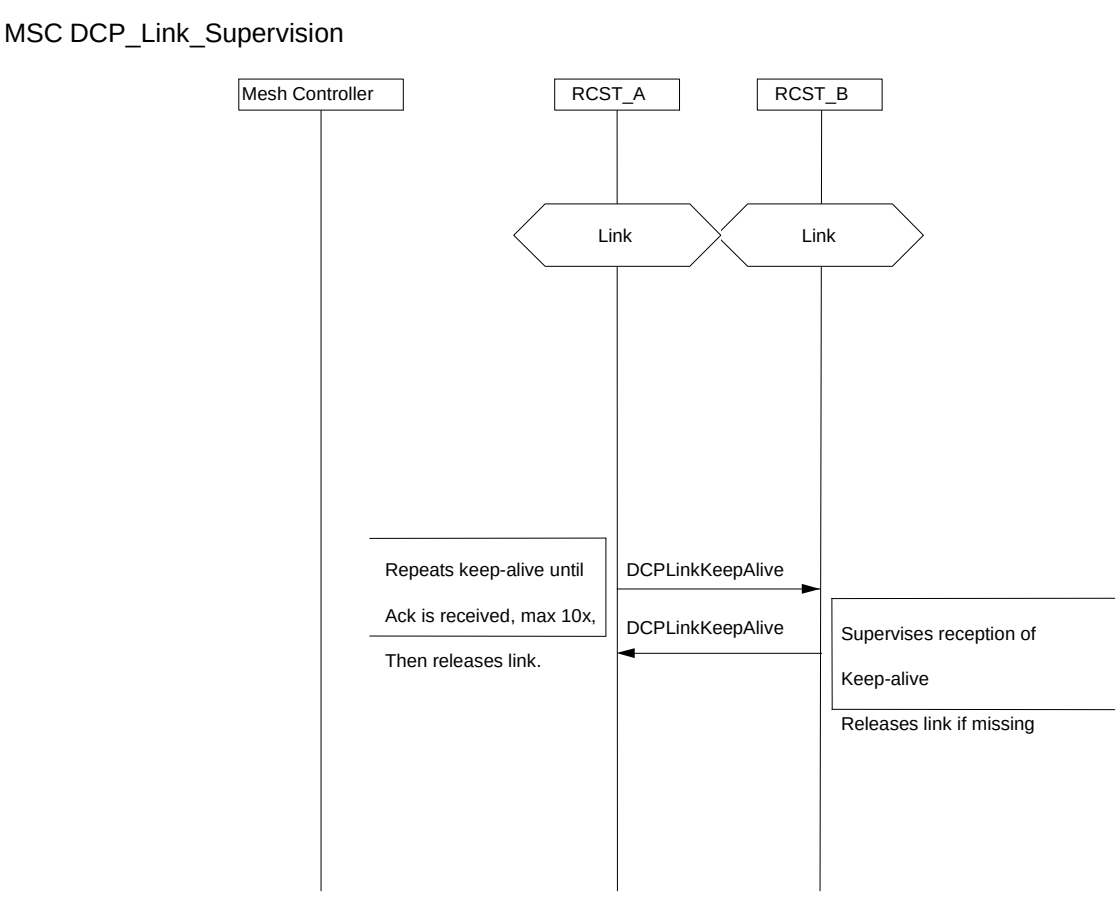


**Figure 15.5: Example of DCP message exchange during link supervision**

## 15.6.4 Link Service Release

A RCST that takes down a link service should do this by requesting the MC to take down both this link service and the corresponding link service used for the opposite direction. The RCST will wait for the response and acknowledge the response, and will release mesh if there is no response from the MC. The other RCST will acknowledge the link service release. The MC will release either RCST from mesh if there is not acknowledgement from the respective RCST.

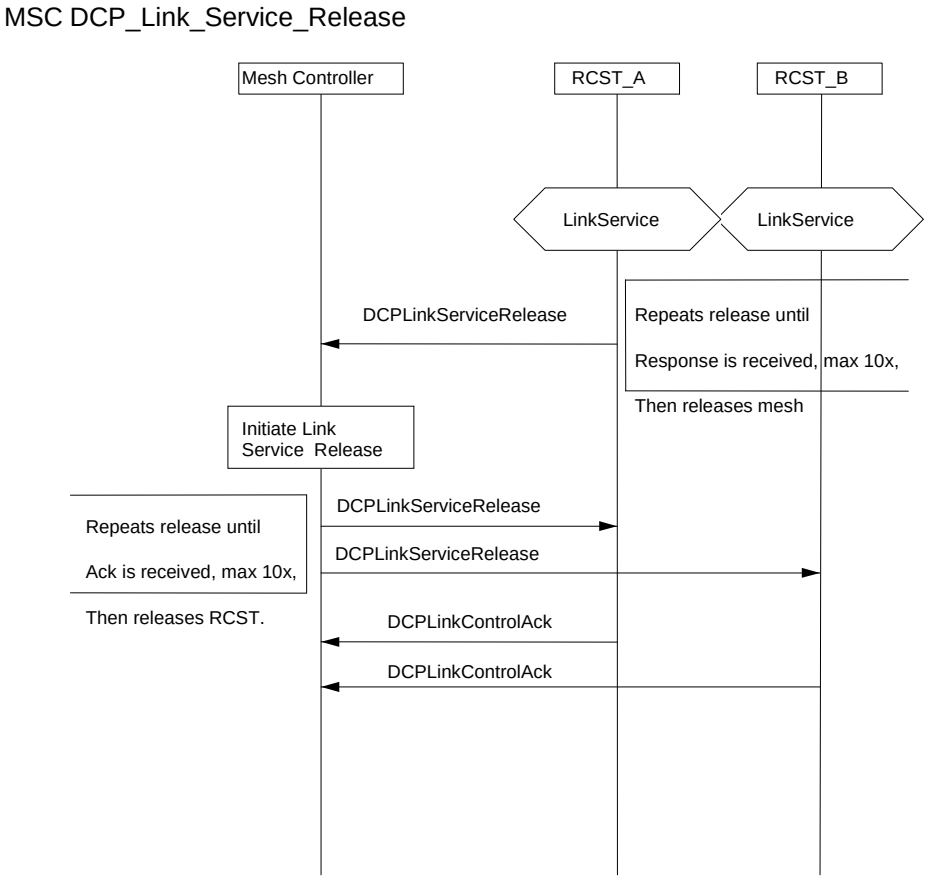


**Figure 15.6: Example of DCP message exchange during link service release**

## 15.6.5 Link Service Keep Alive

A RCST operating a TX link service sends keep-alive regularly to the MC as long as the link service is in use. The MC should autonomously release the TX link service if it does not receive these keep-alive messages in time. The RCST should send keep-alive to the MC at an interval less than half of the shortest timeout interval for the link service, independent of the traffic activity.

The MC should maintain a supervisory timer that autonomously initiates release of a link service if it is not kept alive by the RCST. The MC should simultaneously release the pair of link services used to serve a two-way mesh connection.

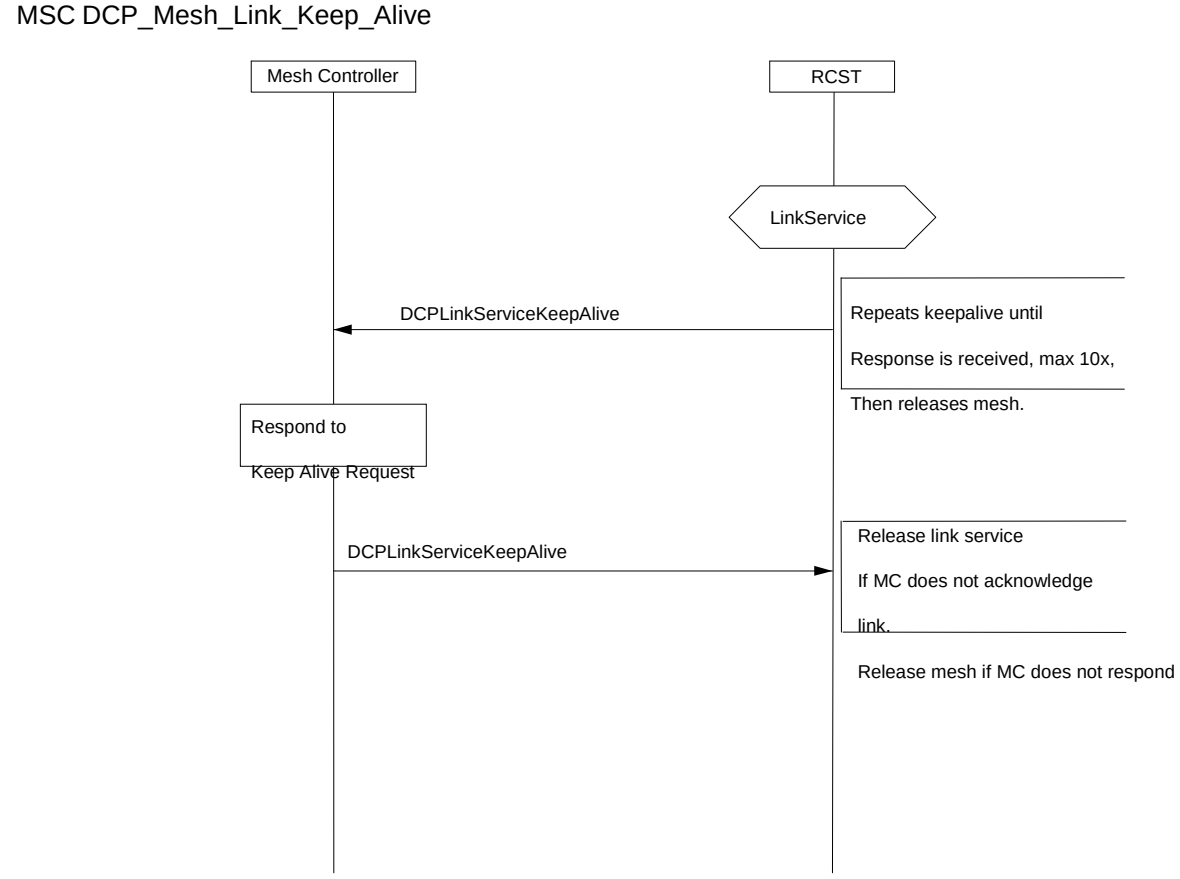


**Figure 15.7: Example of Link Keep Alive**

The RCST maintains one keep-alive timer per TX link service, and issues keep alive when the timer expires. Reception of MC feedback for the specific link service should reset the link specific timer.

# Annex A:
# Interworking with the NGN service layer

## A.1 Policy and Charging Control (PCC) Architecture

Despite the fact that the NGN design is focused to integrate different access networks, satellite networks have many singular characteristics that require further analysis and specification. This clause provides recommendations in regards to the integration of DVB-RCS2 satellite access networks with the Service Layer.

Next Generation Networks (NGN) are all-IP based where all services are supported over IP-backbone (see Figure A.1). A DVB-RCS2 network is viewed in this clause as an example access network that can be integrated in the NGN architecture. In NGN, traditional circuit switched services (e.g. voice) and more feature rich multimedia service are supported over IP-backbone. One main aspect of NGN networks is a clear separation between the lower transport network layers and the upper service layers. This enables a common service environment to be used across different access networks. IP Multimedia Subsystem (IMS) is such a common service environment that is well specified in 3GPP [i.45] and has wide industry support. IMS provides the service requirements that are used by elements in the Policy and Charging Control (PCC) architecture to control the bearers in the underlying access network. Simply stated; IMS is the service layer, PCC is the control and DVB-RCS2 is the transport layer.

The supported IMS services can for example be voice telephony services, video on demand, interactive IPTV, or video surveillance. New services can be developed with available IMS tooling. IP multimedia applications are, as a principle, not standardized, allowing rapid service creation and deployment using standard service capabilities.

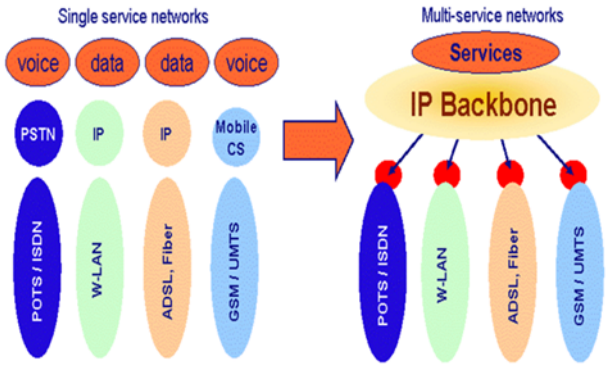


**Figure A.1: Next Generation telecommunication networks are moving away from stove pipe architectures to multi-access / multi-service networks**

Policy and charging control (PCC) rules can be derived using information/requirements provided by the application function (AF). The AF represents applications that require dynamic policy and QoS control in the access network. IMS-based applications provide policy information to the AF. If for example an IMS-multimedia session need to be set-up and maintained, the AF (see Figure A.2) will feed the necessary policy information to the PCC elements (PCRF/PCEF). This clause describes the standard interfaces between the network transport layer and the service control layer as defined in the 3GPP specifications.

The PCC (Policy and Charging Control) architecture is specified by 3GPP [i.46], the architecture specifies both 3GPP access (e.g. UMTS, LTE) as well as non-3GPP (Wifi, Wimax) access networks. A DVB-RCS2 satellite network is viewed as a non-3GPP access network that may use the PCC architecture.

Within the PCC architecture it is possible to set up and control a session with multiple media streams. For many of these media streams, a specific QoS may be required. Especially in a capacity constrained access network like Satellite it is important that for services like voice communications, sufficient bandwidth is reserved and guaranteed. The PCC architecture enables the necessary control of underlying bearers. IMS uses the PCC architecture to control the QoS of bearers and / or IP flows in the IP access network. Figure A.2 depicts the PCC architecture.

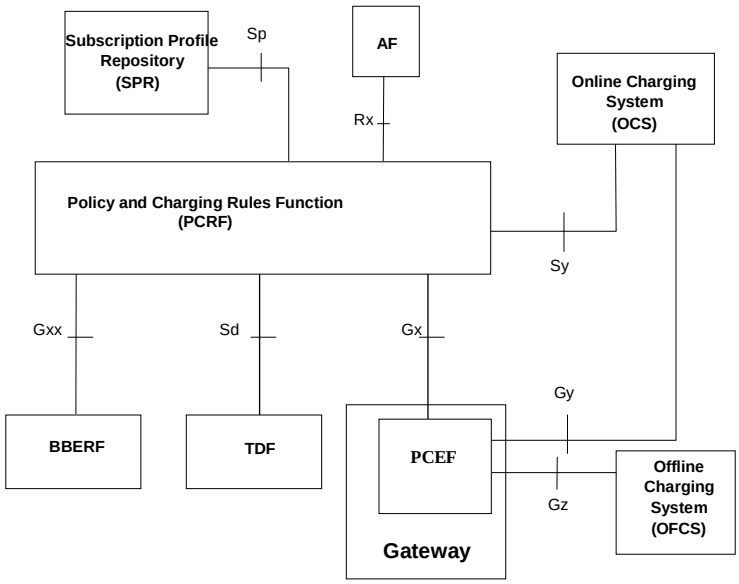


NOTE: PCEF function in the network that enforces policies rules set in the PCRF to control services supported in the underlying all-IP network [i.46].

**Figure A.2: PCC architecture**

The Policy Control architecture consists of the following functional components related to the integration with DVB-RCS2 access network; see Figures A.2 and A.3 where the components and interfaces are detailed:

- Policy Control and Charging Rules Function (PCRF). The PCRF enables the provisioning of policy decisions to policy enforcement functions using PCC/QoS rules. The PCRF performs policy rule authorization for each policy request and assigns a QoS class (QCI) and QoS parameters to each rule for prioritization.

- Policy and Charging Enforcement Function (PCEF). The PCEF has the capability of policing packet flow into an IP network or other transport network (e.g. by controlling a network router, and GWs)

- 3GPP also defines a functional element to control non-3GPP accesses and serving Packet GWs through Gxx interface (Gxa, Gxb or Gxc): "Bearer Binding and Event Reporting Function" (BBERF). BBERF maps the policy decisions from the PCRF to access network specific parameters.

- TDF (Traffic Detection Function) is defined in 3GPP Release 11. The TDF is a functional entity that performs application detection and reporting of detected application and its service data flow description to the PCRF. Using Sd interface, the PCRF may instruct the TDF on which applications to detect and report to the PCRF by activating the appropriate ADC (Application Detection Control) rules. The TDF may be also pre-configured on which applications to detect and report. It is very useful to provide policy control to traffic that is not based on IMS signalling. The TDF to PCRF reference point, listed as Sd (see Figure A.2), have strong similarities to the 3GPP system specific Gx reference point, because the Sd is a subset of the Gx. TDF was introduced to allow implementations where the traffic detection and enforcement functions are separated (e.g. from different vendors). Note that a significant number of operators implement a single vendor Packet Gateway (PDN gateway) that includes the PCEF function and traffic detection function.

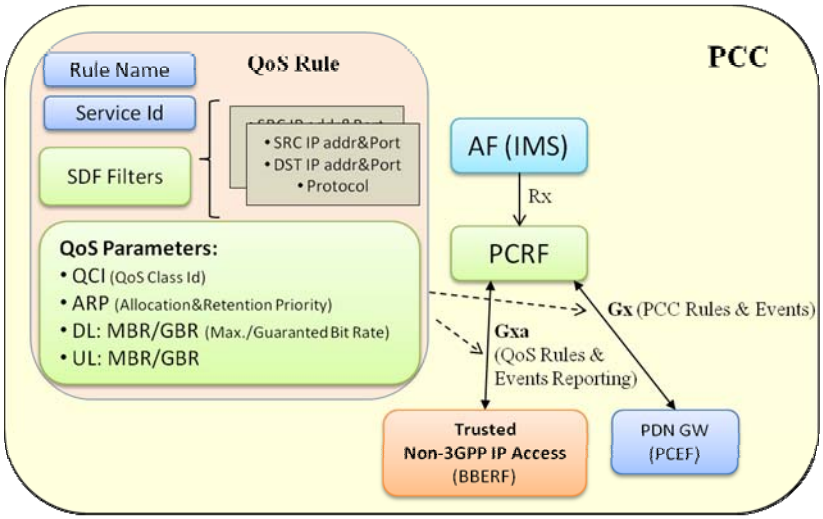


**Figure A.3: QoS Rules in PCC architecture**

The second function in the PCC acronym, after Policy, is Charging Control. Charging the users for services is important in order to support the satellite operator's business. Satellite network operators may benefit from available BSS (Business Support Systems) and already developed systems in the mobile and wireline domain. The available PCC architecture standards, with standard interfaces Gy (for online charging and usage monitoring), or Gz (for off-line charging) should be used to interface with billing systems.

For the purpose of charging correlation between application level (e.g. IMS) and service data flow level, applicable charging identifiers should be passed along within the PCC architecture. The operator should be able to off-line or online charge the users, via standards interfaces to BSS.

The GSM/UMTS core network-charging architecture and principles are specified in [i.47], which provides an umbrella for other charging management documents that specify:

- The content of the CDRs per domain and subsystem (offline charging).
- The content of real-time charging messages per domain / subsystem (online charging).
- The functionality of online and offline charging for those domains and subsystems.
- The interfaces that are used in the charging framework to transfer the charging information (i.e. CDRs or charging events).

Subscription information contained in the SPR (Subscription Profile Repository) or HSS (Home Subscriber Server), see Figure A.2, is used to set policy rules for a particular user. For example the QoS subscription information may be used to derive a policy rule that is used to enforce the maximum data rate of a service data flow that the user has subscribed to.

## A.2 Integrating DVB-RCS2 Access Network into the PCC architecture

Figure A.4 shows a generic multi-service network scenario where the UE (User Equipment) is using a satellite network (OVN defined in HLS document) to access to different services. In the figure some UEs are subscribed to the Internet service provided by the Internet Service Provider (ISP), other UEs are subscribed to a wholesale ISP, and others to 3GPP/NGN/IMS services. The UEs are attached to "Customer Premises Networks (CPN)", typically Ethernet LANs. And each CPN uses a "Customer Network Gateway (CNG)" as access device.

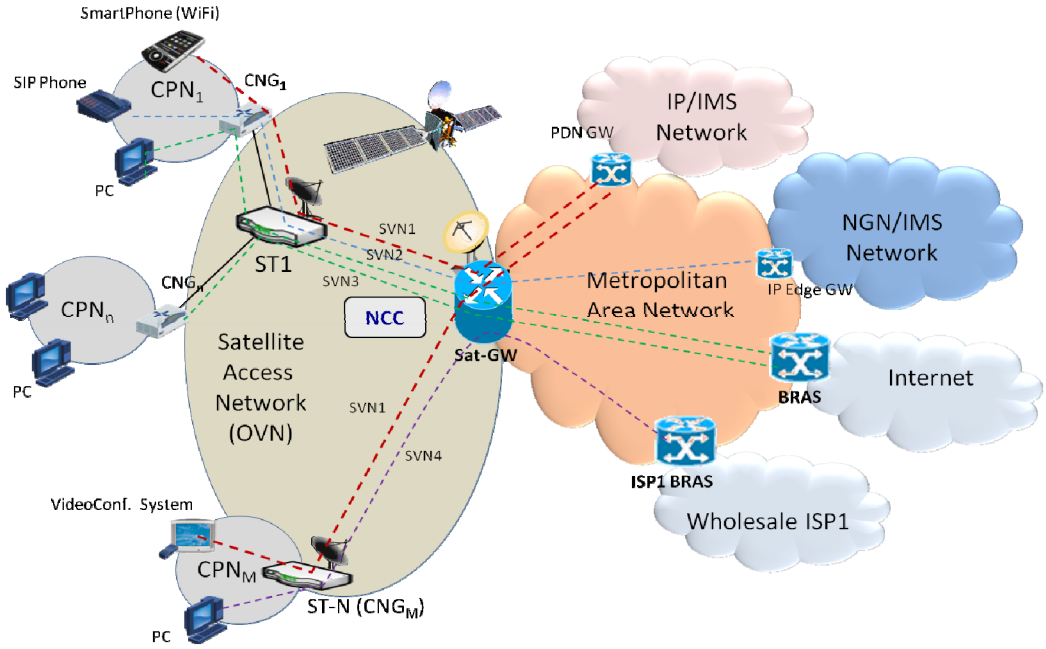


**Figure A.4: General multi-service Satellite Access Scenario**

Several CNGs can be connected to a satellite terminal (ST – e.g. DVB-RCS2 RCST). The satellite terminal may function as the CNG itself. Different SVNs (Satellite Virtual Networks) may be used to segregate the traffic associated to different services. In Figure A.4, each SVN (with an individual colour) is one-to-one mapped to end service provider. In this scenario, the role of the OVN in the complete PCC architecture may be different depending on the integration approach. Two different approaches have been identified:

1) PDN GW (Packet Data Network GateWay) integrated in the OVN (Sat GW/NCC): This is a very satellite-centric approach where OVN is the only access network for users. In this approach, the OVN implements the PCEF and TDF functional components, interacting with the rest or the PCC components through Gx, Gy, Gz and Sd interfaces. In this approach, OVN implements both policy and charging control. Figure A.5 highlights the components and control interfaces of the OVN with this approach.

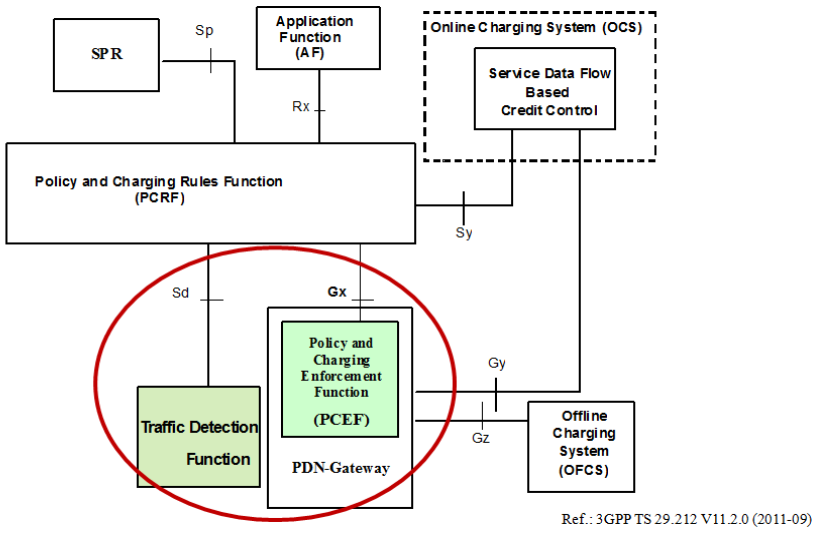


**Figure A.5: PDN GW integrated in the OVN (Sat GW)**

2) OVN as a trusted access network of a general multi-service/multi-carrier network: In this approach, the OVN shares most of the PCC components with other access networks and service providers. This second approach is much simpler in terms of functionality and interfaces; and it is fully specified in PCC standards [i.48]. Also, it provides a real integration of the OVNs with the already existing networks. Figure A.6 highlights the components and control interfaces of the OVN with this approach.

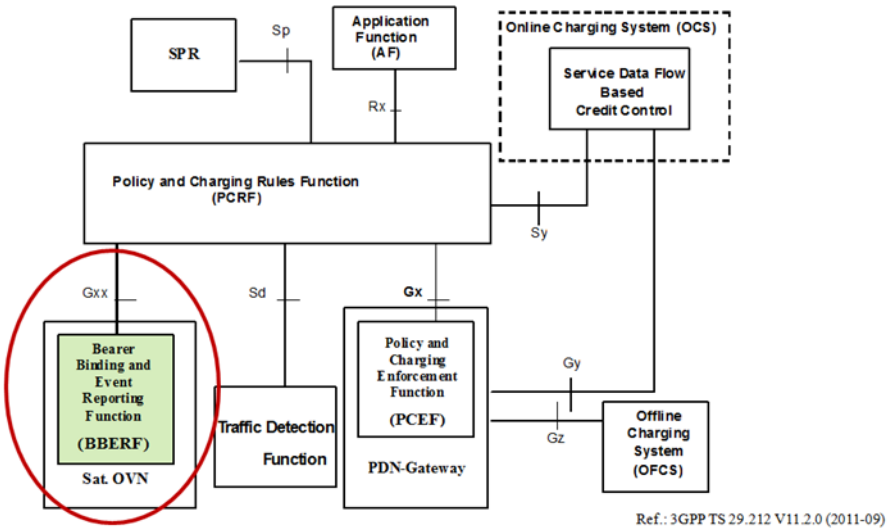


**Figure A.6: OVN as a trusted access network**

For this second approach, different integration schemes have been defined in [i.48] for trusted IP access networks; all of them use Gxx (Gxa) as control interface (QoS rules provisioning and Event reporting) and S2a or S2c as data interfaces. In this guideline document the simplest scheme has been selected; it provides:

- Full QoS control (through Gxa) for IMS sessions traffic and also for non-IMS traffic detected by TDF. In this scheme, the DVB-RCS2 network (OVN) should implement the BBERF functionality.
- User mobility is based on "PMIPv6 Network Mobility" where UEs do not need any modification to support mobility. DVB-RCS2 network should implement the S2a PMIPv6 interface defined in [i.48].

Figure A.7 highlights the data and control interfaces of the OVN behaving as a trusted Non-3GPP IP Access.

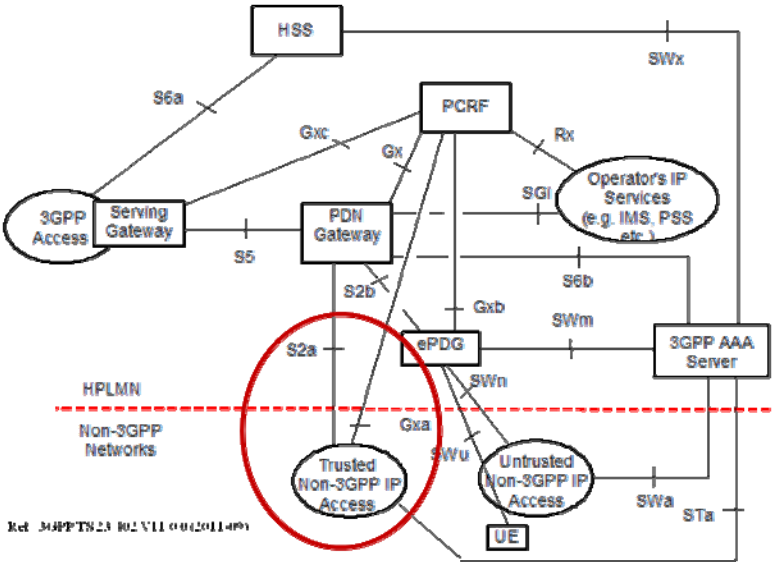


**Figure A.7: Data and Control interfaces for OVN as Trusted Access**

# A.3 Interfaces and Reference Points

The Reference Points in Table A.1 (which have the properties of interfaces) are those between DVB-RCS2 network entities and NGN entities within a DVB-RCS2 System and assume that certain entities are integrated into DVB-RCS2 entities (e.g. the BBERF/PCEF) as indicated in previous clauses. The interfaces are all based on the 3GPP definitions for these reference points.

**Table A.1**

| Reference Point | Entities | Use |
|---|---|---|
| Gx (Diameter) | PCEF - PCRF | Policy enforcement and control |
| Gxa (Diameter) | PCRF and the BBERF | Policy enforcement and control [i.46] |
| Gy (Diameter) | PCEF – OCS (online charging system) | Online charging, online usage meeting for gating and/or throttling |
| Gz | PCEF – OFCS (Offline Charging System) | Call data records (CDRs) FTP file transfer |
| Ro (Diameter) | CSCF – OCS | Used to exchange online charging information with OCS |
| Rx (Diameter) | PCRF – CSCF | Used to exchange policy and charging related information between P-CSCF and PCRF |
| ISC (SIP) | CSCF – AF | Notify the AF of registration state, UE capabilities, etc. |
| Sp | HSS or SPR | provide subscription data to PCC |
| Sd | TDF – PCRF | Policy and Charging control |
| S2a | 3GPP-PDN – Trusted Non-3GPP-PDN | 3GPP interface to Trusted Non-3GPP IP access network. It supports of mobility management of mobile devices [i.48] |

# A.4 Interactions with DVB-RCS2 network

## A.4.1 Interaction between the PCEF/BBERF and PCRF

The interface between the PCRF and PCEF/BBERF is via standard Diameter based interfaces. The policies that are defined in the PCRF are sent to the PCEF over the Diameter Interface. Note that it is not necessary to manage all session via the PCEF/BBERF, this depends on the enforcement rules set on the PCEF/BBERF. The policies are defined in AVPs (Attribute Value Pairs) that are exchanged between the PCRF and PCEF/BBERF. The full set of AVPs is specified in [i.47] and [i.49]. This is a standardized interface allowing different PCRF systems from different vendors and operators to connect over this interface. This is common practice in mobile networks where different service providers share the same network and each Service Provider (SP) is able to set policies (within an agreed set) for its own subscribers.

PCRF provisions QoS rules to BBERF component (or PCEF if the first approach previously defined is used, where QoS rule is contained into a more general PCC rule) implemented into the OVN.

## A.4.2 Mapping of BBERF/PCEF to DVB-RCS2 controls

For the policies to be enforced through the BBERF/PCEF function there should be a mapping that is DVB-RCS2 internal. This mapping should ensure that – for example – a required guaranteed bit rate for a voice service is enforced.

The BBERF/PCEF has to configure all the satellite components involved in the IP flows contained in the QoS rule as SDF filters (Service Data Flows, See Figure A.3); both forward and return flows: NCC, GW and ST. To perform this configuration, new control interfaces needs to be defined internally in the satellite network as Figure A.8 shows. A SDF is an aggregate set of packet flows that matches a service data flow template (IPs, ports, etc.).

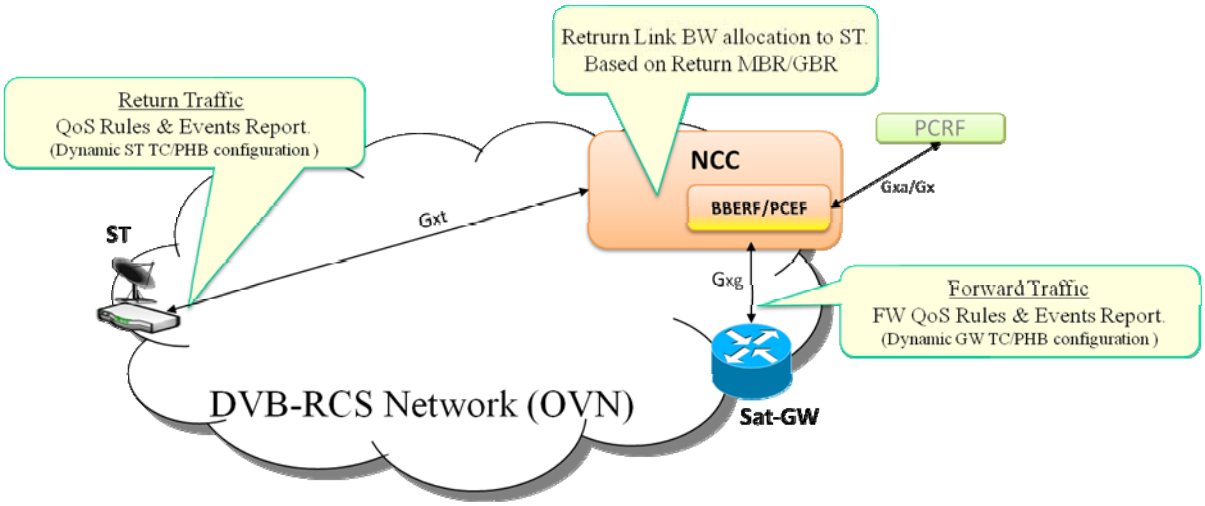


**Figure A.8: DVB-RCS2 network Internal Policy Control interfaces**

First, BBERF/PCEF should bind the UE´s IP address to the ST it is sending traffic through. Afterwards, NCC needs to allocate enough return bandwidth to the ST based on the Return Maximum & Guarantied Bit Rates (MBR/GBR), and interact with the ST and GW to provision the return and forward QoS Rule. The terminal and the GW should use the rule to dynamically configure the Traffic Classification and Per Hop Behaviours for the corresponding IP flows.

These two new internal control interfaces (called in the present document as Gxt and Gxg) need to be fully defined. The proposal for these new interfaces should be based on the "all IP" interfaces already defined for Policy Control, such as Gxx. Both Gxt and Gxg interfaces provision QoS rules according to [i.49] using diameter AVPs over TCP or SCTP connections between NCC and STs/GW. A more detailed definition of these new interfaces is provided below.

## A.4.3 Policy control on the RCST&GW

Many UEs can be connected to a RCST, via for example a LAN interface. When an UE wants to setup a session with a particular QoS, the RCST will need to be aware of the QoS requirements for the return traffic of the requested session (e.g. Video Call).

The "current approach" to control the traffic in a RCST is shown in Figure 7.4 in clause 7 where user plane and satellite bearer control functions are completely separated. The RCST classifies and schedules packets using a static configuration (e.g. managed through SNMP). It also controls the satellite resources using BoD techniques based on "Traffic Snooping" and "Buffers Monitoring". Because the RCSTs are unaware of service logic and PCC protocols, they should determine QoS parameters based on static configured rules and traffic snooping.

Figure A.9 shows a new enforcement module (T-BBERF) and interfaces in the RCST as an extension to the "current approach" in Figure 7.5. T-BBERF module uses the QoS rules and the current available channel capacity to dynamically configure the classification, shaping/policing and packet scheduling (in general, Traffic Classification & Per Hop Behaviour, TC&PHB). A binding between the QCI contained in the QoS rule and a DiffServ Class should be configured and applied in the T-BBERF.

Also, T-BBERF may inform the NCC about the status of the different aggregated flows per SVN and quality class as event reports. NCC capacity assignment can take into account both, the QoS Rules and the status reports (note that status reports are also useful for non policy controlled traffic).

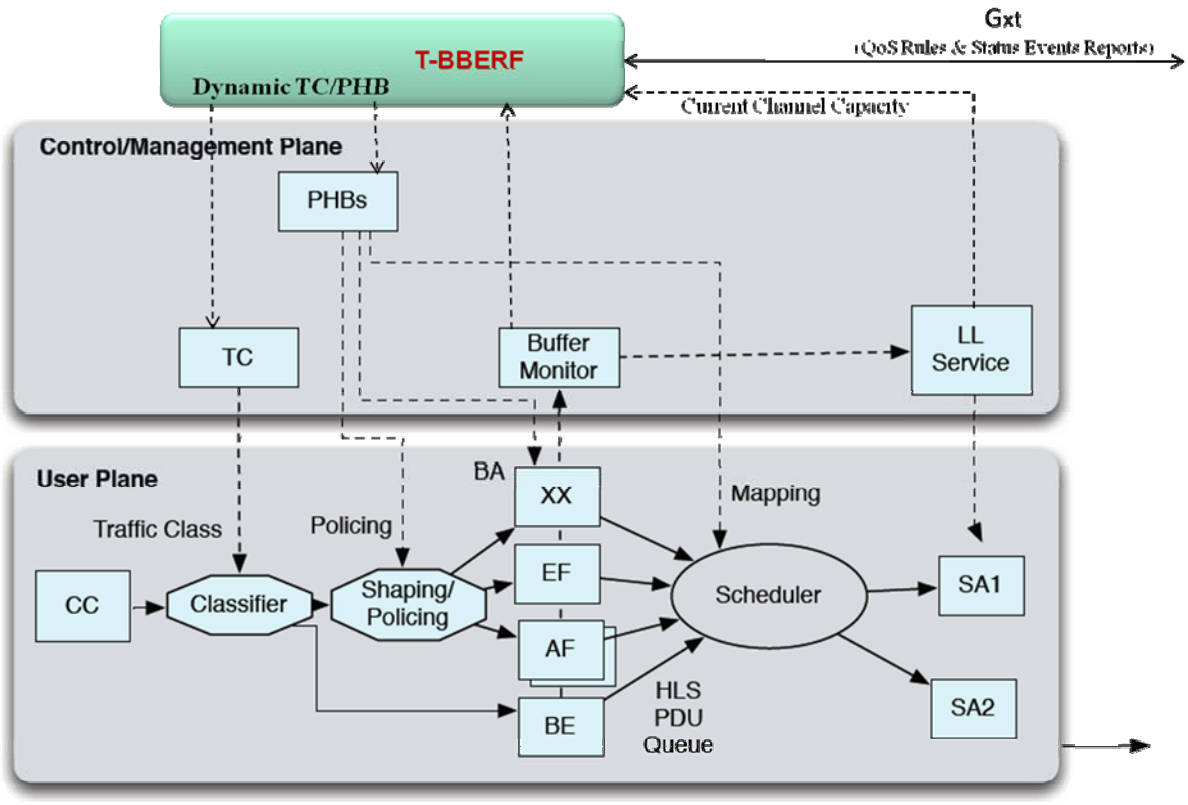


**Figure A.9: New T-BBERF Component in RCST Control Plane**

The internal behaviour of the new T-BBERF component is out of the scope of the guideline document; different implementations are possible and they do not affect the ST interoperability if it is compliant with the Gxt interface defined below.

The same functional QoS enforcer module (called GW-BBERF in this case) should be implemented into the satellite GW to control forward traffic QoS.

Both T-BBERF and GW-BBERF provide dynamic control over the user plane traffic handling and encompasses the functionalities defined in [i.48], section 4a.4.2 for the BBERF component. These functionalities are mainly:

- It should ensure that the service data flow under QoS control is carried over the return or forward satellite bearer with the appropriate QoS class. The ARP, GBR, MBR and QCI parameters in the QoS Rules (see Figure A.3) are used for selecting the appropriate PHB (e.g. Weights of the packet scheduler).
- Event reporting: It should report events to the NCC based on the event triggers installed by the NCC using the Gxt/Gxg procedures defined below.

## A.5 Example of a SIP call

Figure A.10 shows an example of the steps required to complete a SIP call (e.g. Video Call) when the DVB-RCS2 Satellite Network (IP-CAN: IP Connectivity Access Network. A general term used to denote an Access Network that provides IP connectivity) is integrated with the PCC architecture.

We assume that the BBERF has already established a Gateway control session with the PCRF as specified in [i.49]. Detailed control sequences can be found in [i.50]. These steps can be summarized as follows:

1) An UE requests a session (SIP call) to the "Call Session Control Function (CSCF)" of the IMS system acting as "Application Function (AF)".

2) The CSCF use the Rx reference point to exchange application level session information with the Policy and Charging Rules Function (PCRF). This information is part of the input used by the PCRF for the Policy and Charging Control (PCC) decisions.

3) After admission control, the PCRF generate the corresponding QoS rule. It uses the standard Gxa interface to provision the QoS rules in the BBERF implemented in the NCC. This rule is bind to the Satellite IP-CAN connection (one per ST and SVN).

4) The rule enforcement required of the two new control interfaces internal to the Satellite Network (Gxt and Gxg) to enforce the rule in both the Sat-GW and the ST; where the enforcement functions should be implemented as discussed above to provide dynamic TC/PHB configuration.

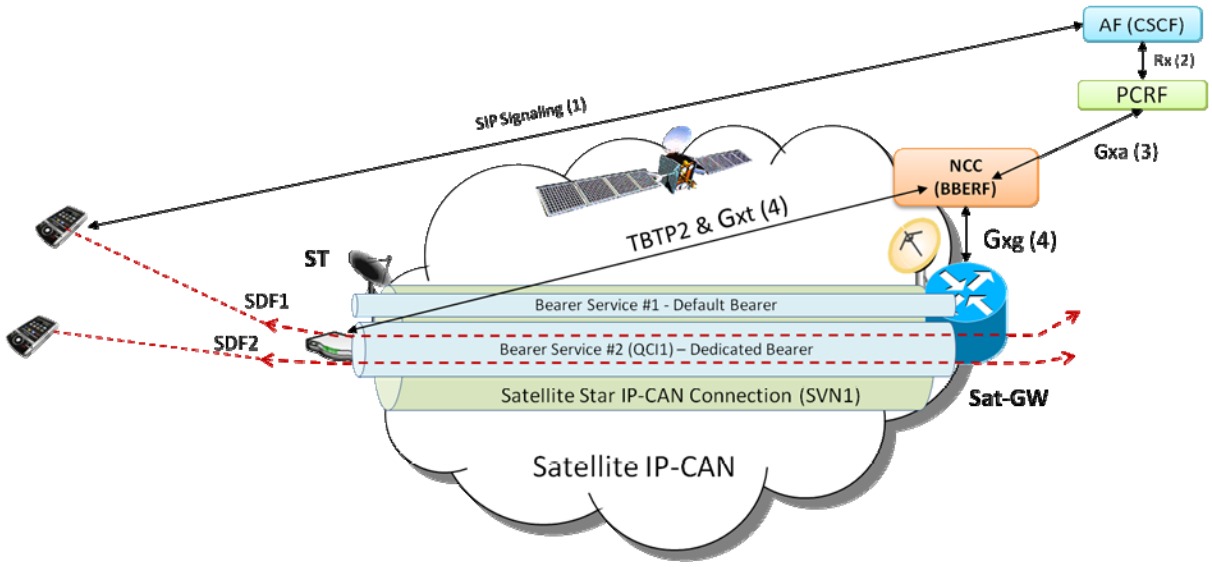


**Figure A.10: Example, Steps complete a SIP Call with PCC**

The concrete implementation of an IP-CAN connection and a Bearer Service in the satellite network can be implementation dependent. As an example, the IP-CAN connection can be the portion of a DVB-RCS2 "connectivity channel" used by a SVN, and the Bearer Services can be the DiffServ Class associated to the QCI.

Next clause details the signalling flows involved in this example; see Figures 14.16 and A.17.

# A.6 Gxt and Gxg Reference Points

The Gxt reference point is located between the NCC and the T-BBERF (Satellite Terminal Bearer Binding and Event Reporting Function). The Gxg reference point is located between the NCC and the GW-BBERF (Satellite Gateway Bearer Binding and Event Reporting Function).

The Gxt and Gxg reference points are used for:

- Provisioning, update and removal of QoS rules from the NCC to the T/GW-BBERF.
- Transmission of traffic plane events from the T/GW-BBERF to the NCC.

These reference points are proposed to be fully compliant with the Gxx reference point defined in [i.49], (section 4a) where the NCC has the functionality of Policy Controller (PCRF) and the T/GW-BBERF has the functionality of BBERF.

We provide below a description of the procedures and signalling flows involved in the policy control in a DVB-RCS2 network. The procedures and signalling flows for session termination and session modification are not provided in the present document, but all of them are compliant with Gxx procedures and protocol defined in [i.49].

Alternatively to the use of Diameter sessions for the Gxt/Gxg reference points, the provisioning of QoS rules and status reporting could be carried out using DCP. Clause 13.2.5 of the present document specifies other possible DCP functionalities, including dynamic QoS provisioning. Some adaptation to the already defined IEs, new IEs or new DCP messages may be necessary to include dynamic traffic classification rules. This adaptation should be possible since DCP has been specified to allow this degree of flexibility.

## A.6.1 Initial Satellite Terminal and Gateway Attachment procedure

When a Satellite Gateway or Terminal starts-up it should establish a Diameter connection with the NCC. This connection will be used to send and receive all the Diameter messages related to Policy Control. Document [i.49], section 5a.2 details the Gxx procedures of "Initialization, maintenance and termination of connection and session" that we apply for Gxt and Gxg interfaces.

With regard to the Diameter protocol defined over the Gxt/Gxg interface, the NCC acts as a Diameter server.
The T-BBERF or GW-BBERF acts as the Diameter client.

Figure A.11 shows the signalling flow required when the Sat-GW starts-up.

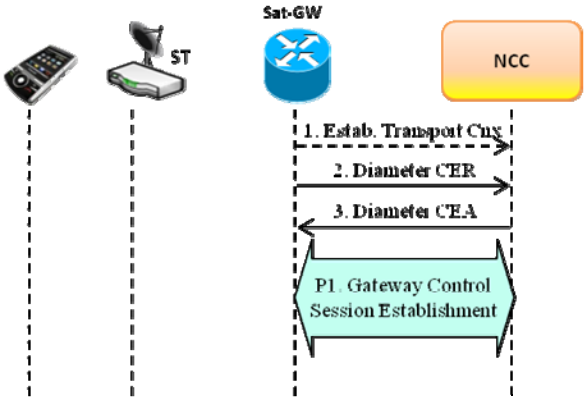


**Figure A.11: Initial Satellite Gateway Attachment**

Where:

1. The GW-BBERF (Sat-GW) establishes the transport connection. The initialization of the connection between the GW-BBERF (Sat-GW) and NCC is defined by the underlying transport protocol: TCP port 3864.
2. and 3. After establishing the transport connection, the NCC and the GW-BBERF should advertise the support of the Gxg specific Application using the CER (Capabilities Exchange-Request) and CEA (Capabilities-Exchange-Answer) commands specified in the Diameter Base Protocol [i.51].
"P1. Gateway Control Session Establishment": The GW-BBERF initiates this procedure in order initiate the policy control with the NCC; in this session establishment the default QoS rules and event triggers for all the SVNs that the Sat-GW handles may be deployed. This procedure "P1. Gateway Control Session Establishment" is defined in the next clause.

Figure A.12 shows the signalling flow required when a ST starts-up.

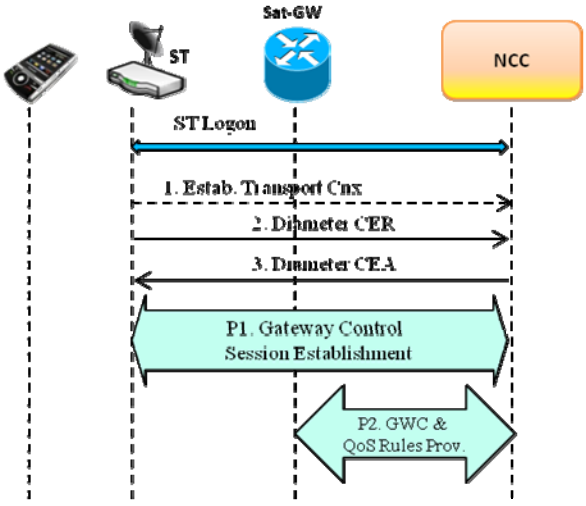


**Figure A.12: Initial Satellite Terminal Attachment**

Where:
1. After ST logon, the T-BBERF (ST) establishes the underlying TCP connection.
2. and 3. After establishing the transport connection, the NCC and the T-BBERF should advertise the support of the Gxt specific Application using the CER and CEA.
"P1. Gateway Control Session Establishment": The T-BBERF initiates this procedure in order initiate the policy control with the NCC; in this session establishment the NCC may deploy the default QoS rules and event triggers for all the SVNs that the ST handles. This procedure "P1. Gateway Control Session Establishment" is defined in the next clause.
"P2. Gateway Control & QoS Rules provision": The NCC deploys in the Sat-GW the default QoS rules and event triggers required for the traffic with the ST. This procedure "P2. Gateway Control & QoS Rules provision" is defined in the next clause.

The procedures P1 and/or P2 are used in most of the policy control signalling flows, and they are defined in a general way in the next clause.

## A.6.2 Gateway Control Session Establishment Procedure on Gxa, Gxt and Gxg

The Gateway Control Session Establishment Procedure on Gxx interface is fully defined in [i.49], section 4a.5.1 and [i.50], section 4.4.1. Note that the procedure has been simplified, not including the roaming scenarios fully defined in [i.49]. Figure A.13 shows this procedure that can be also applied on Gxt and Gxg interfaces:

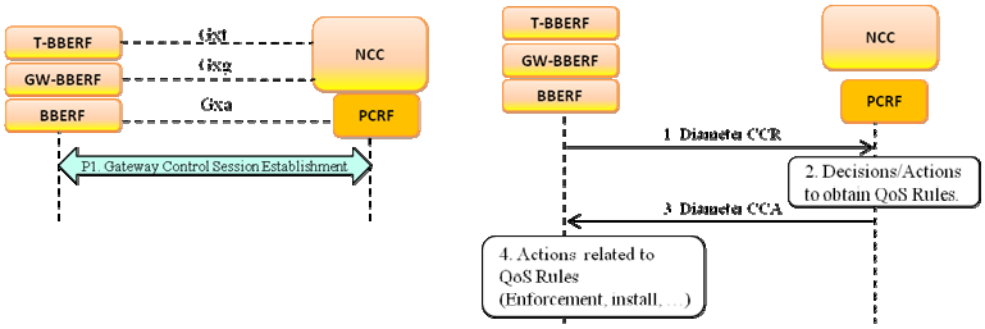


**Figure A.13: P1. Gateway Control Session Establishment**

1) The BBERF initiates a Gateway Control session with the PCRF by sending a CCR to the PCRF with the CC-Request-Type AVP set to the value INITIAL_REQUEST. The BBERF provides equipment identity and other information as defined in [i.49]. For the new T-BBERF component, the equipment identity may be the logon_Id of the ST when it is attached to the network; or it may be the UE MAC address when the UE is attached. The mapping between this identities and the IMSI subscriber identification required in Gxa interface should be performed by the NCC based on configuration data.

2) The NCC or PCRF performs the following actions:

   - It stores the information received in the CCR.
   - If it requires subscription-related information and does not have it, it requests such information.
   - It prepares for the installation of QoS rules if available.
   - It stores the selected QoS Rules and PCC Rules.

   On UE attachment, the NCC stores the binding of the Gxt session with the associated Gxa session.

   The PCRF may correlate the UE identity information with already established Gx sessions for the same UE.

3) The NCC or PCRF acknowledges the Gateway Control Session by sending a CCA to the T/GW-BBERF. It includes the available QoS rules and the event triggers.

4) The T/GW-BBERF installs and enforces the received QoS Rules.

## A.6.3 Gateway Control & QoS Rules Provision Procedure on Gxa, Gxt and Gxg

The Gateway Control & QoS Rules Provision on Gxx interface is fully defined in [i.49], section 4a.5.2 and [i.50], section 4.4.3. Note that the procedure has been simplified, not including the roaming scenarios fully defined in [i.49]. Figure A.14 shows this procedure that can be also applied on Gxt and Gxg interfaces:

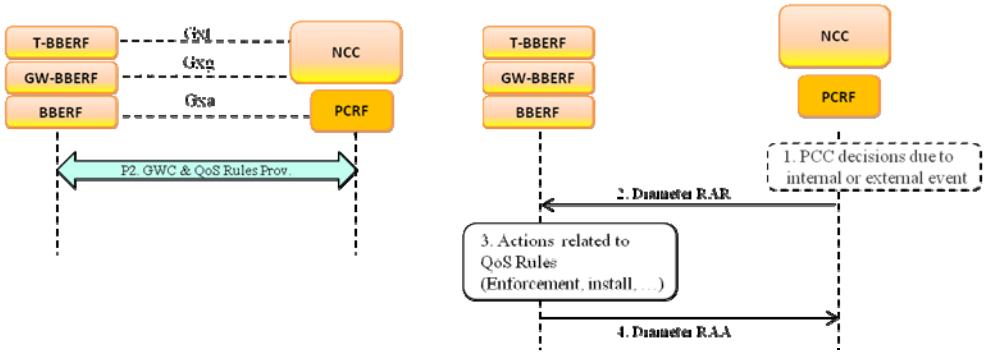


**Figure A.14: P2. Gateway control and QoS Rules Provision**

1) The NCC or PCRF receives an internal or external trigger to update QoS Rules and event triggers for a gateway control session. The NCC/PCRF may decide to operate on QoS Rules without obtaining a request from the T/GW-BBERF, e.g. in response to information provided to the NCC via the Gxa reference point, or in response to an internal trigger within the NCC/PCRF.

2) The NCC or PCRF sends a Diameter RA-Request message (RAR) to request that the T/GW-BBERF installs, modifies or removes QoS Rules and/or updates the event triggers.

3) The T/GW-BBERF installs, modifies or removes the identified QoS Rules. The T/GW-BBERF also enforces the authorized QoS and enables or disables service flow according to the flow status of the corresponding QoS Rules.

4) The BBERF sends a Diameter RA-Answer message (RAA) to the NCC/PCRF to acknowledge the RAR and informs it about the outcome of the QoS rule operation. If the corresponding resource cannot be established or modified, then the T/GW-BBERF should reject the activation of a QoS rule as specified in [i.49].

## A.6.4 User Equipment (UE) Attachment procedure

There are many possible procedures to complete the attachment of a UE to the DVB-RCS2 network with support of policy control. The selected PMIPv6 S2a [i.48] interface permits that any UE, with no protocol modifications, can be attached to NGN network through using the satellite access, including IP mobility features. Section 4.7.2 of [i.48] defines different UE IP addressing schemes that can be applied, from the static allocation to dynamic allocation based on DHCPv4/v6. Also, section 6.2 of [i.48] defines the ignition attach procedure on S2a interface.

As an example, Figure A.15 summarized the signalling flows of the UE attachment based on DHCPv4 and PMIP compliant with S2a interface:

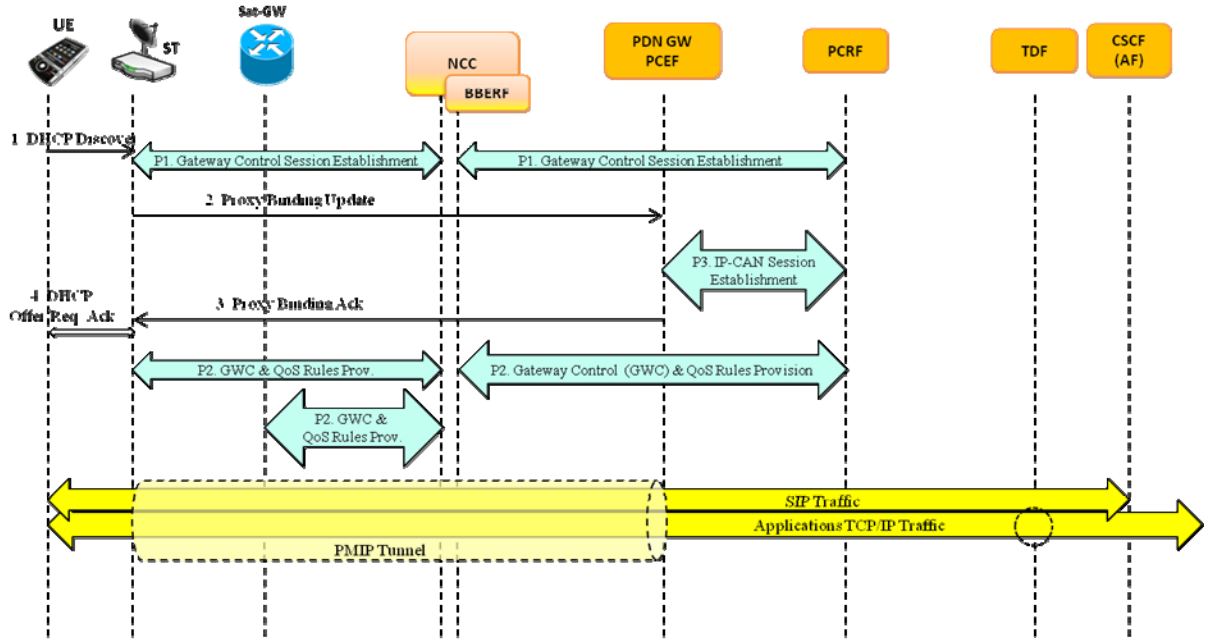


**Figure A.15: Example of UE attachment procedure**

1. The UE sends a DHCPv4 Discovery message in broadcast to the network to find available servers.

"P1: Gateway Control Session Establishment". The ST initiates the Gateway Control Session Establishment Procedure with the NCC and the NCC (BBERF) with the PCRF, as already defined. The DVB-RCS2 access network provides the information to the PCRF to correctly associate it with the IP CAN session to be established in step "P3".

2. Applying the PMIP architecture [i.52], the ST behaves as the MAG (Mobile Access Gateway) and the PDN GW as the LMA (Local Mobility Anchor). The ST sends a Proxy Binding Update (PBU) message to the PDN GW in order to request the new IPv4 address and update the current registration. Upon receiving the PBU message from the ST, the PDN GW allocates an IPv4 address for the UE in accordance with the operator's policies.

"P3: IP CAN Session Establishment". The PDN GW initiates the IP CAN Session Establishment Procedure with the PCRF, as specified in [i.46]. The PDN GW provides information to the PCRF used to identify the session and associate Gateway Control Sessions established in "P1" correctly. The PCRF creates IP CAN session related information and responds to the PDN GW with PCC rules and event triggers.

3. The PDN GW responds with a PMIP Binding Acknowledgement (PBA) message to the ST with the assigned IPv4 Address.

4. The ST acting as a DHCPv4 server sends the DHCPv4 Offer with the assigned UE IPv4 address received in the PBA message in previous step. When the UE receives the lease offer, it sends a DHCPREQUEST message containing the received IPv4 address. The ST sends a DHCPACK packet to the UE. This message includes the lease duration and any other configuration information that the client might have requested.

When receiving the DHCPACK message, the UE completes TCP/IP configuration process.

"P2: GW Control & QoS Rules Provision". The PCRF updates the QoS rules in the DVB-RCS2 access network by initiating the GW Control & QoS Rules Provision Procedure. The NCC also updates the QoS rules in the ST and GW by initiating the GW Control & QoS Rules Provision Procedure.

SIP and Applications traffic can now be sent through the configured PMIP tunnel (e.g. GRE) between ST and PDN GW, using the required service provider addressing plan.

## A.6.5 Signalling flows for IMS

In [i.49] ("Annex B: Signalling Flows for IMS" and "Section 4.3.1: Network-Initiated IP-CAN Session Modification") we can find how IMS signalling is integrated with the PCC procedures that we have integrated with the proposed DVB-RCS2 policy control. Figure A.16 shows the PCC Procedures for IMS Session Establishment at originating CSCF and PCRF, where provisioning of service information is derived from SDP offer and answer.

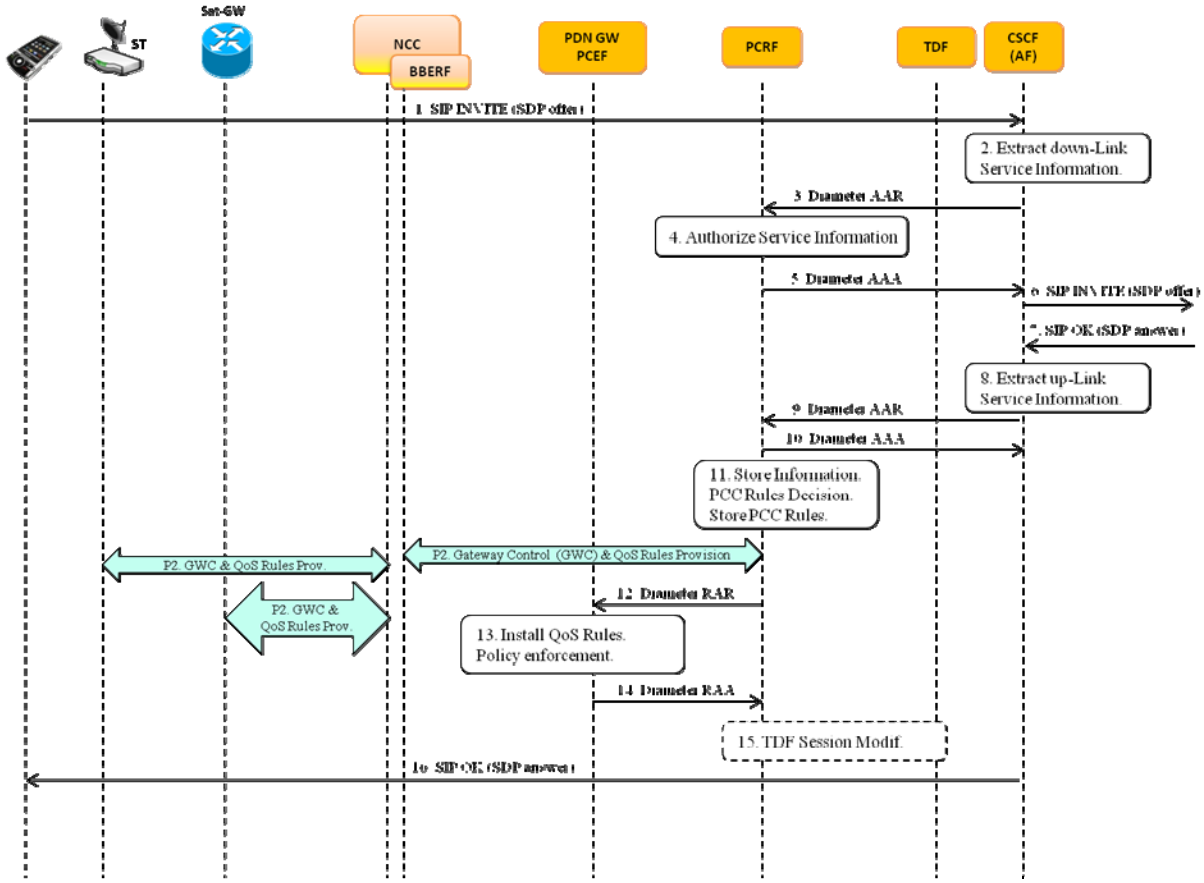


**Figure A.16: Signalling flow for IMS SIP call**

Where:

1. The CSCF receives the first SDP offer for a new SIP dialogue within a SIP INVITE request.

2. The CSCF extracts service information from the SDP offer (IP address of the down link IP flow(s), port numbers to be used etc.).

3. The CSCF forwards the derived service information to the PCRF by sending a Diameter AAR over a new Rx Diameter session. It indicates that only an authorization check of the service information is requested.

4. The PCRF checks and authorizes the service information, but does not provision PCC/QoS rules at this stage.

5. The PCRF replies to the CSCF with a Diameter AAA.

6. The CSCF forwards the SDP offer in SIP signalling.

7. The CSCF receives the negotiated SDP parameters from the terminating side within a SDP answer in SIP signalling.

8. The CSCF extracts service information from the SDP answer (IP address of the up-link media IP flow(s), port numbers to be used etc.).

9. The CSCF forwards the derived service information to the PCRF by sending a Diameter AAR over the existing Rx Diameter session.

10. The PCRF replies to the CSCF with a Diameter AAA.

11. The PCRF selects the PCC Rule(s) to be installed, modified or removed for the IP-CAN Session. The PCRF may also update the policy decision by defining an authorized QoS and enable or disable the service flow(s) of PCC Rules. The PCRF may add or change QoS information per QCI applicable to that session. The PCRF may update the ADC decisions and select the ADC rules to be installed, modified or removed for the session. PCRF stores the updated PCC Rules, and ADC rules.

"P2. Gateway Control & QoS Rule Provision". The PCRF initiates "Gateway Control and QoS rules provisioning procedures" following signalling flows described in Figure A.14.

12. The PCRF sends a Diameter RAR to request that the PCEF installs, modifies or removes PCC Rules and updates the policy decision. In the case of PCEF supporting Application Detection and Control feature, the PCRF may also request the PCEF to install, modify or remove the ADC rules by updating the ADC decisions for the session.

13. The PCEF installs, modifies or removes the identified PCC Rules. The PCEF also enforces the authorized QoS and enables or disables service flow according to the flow status of the corresponding PCC Rules. If QoS information is received per QCI, PCEF should set/update the upper limit for the MBR that the PCEF assigns to the non-GBR bearer for that QCI. In the case of PCEF supporting Application Detection and Control feature, when the solicited application reporting applies, the PCEF may also install, modify or remove the provided ADC Rules.

14. The PCEF sends a Diameter RAA to acknowledge the RAR. The PCEF informs the PCRF about the outcome of the PCC rule operation

15. In case of TDF, solicited application reporting, PCRF initiates the TDF session establishment, modification, or termination.

16. Upon successful authorization of the session, the SDP parameters are passed to the UE in SIP signalling.

Figure A.17 is the same as Figure A.16 but it expands all the signalling flows. From this figure we can obtain the performance impact of the proposed PCC integration scheme. The complete SIP call requires only two additional satellite hops, having a total call establishment delay 1 second approximately.

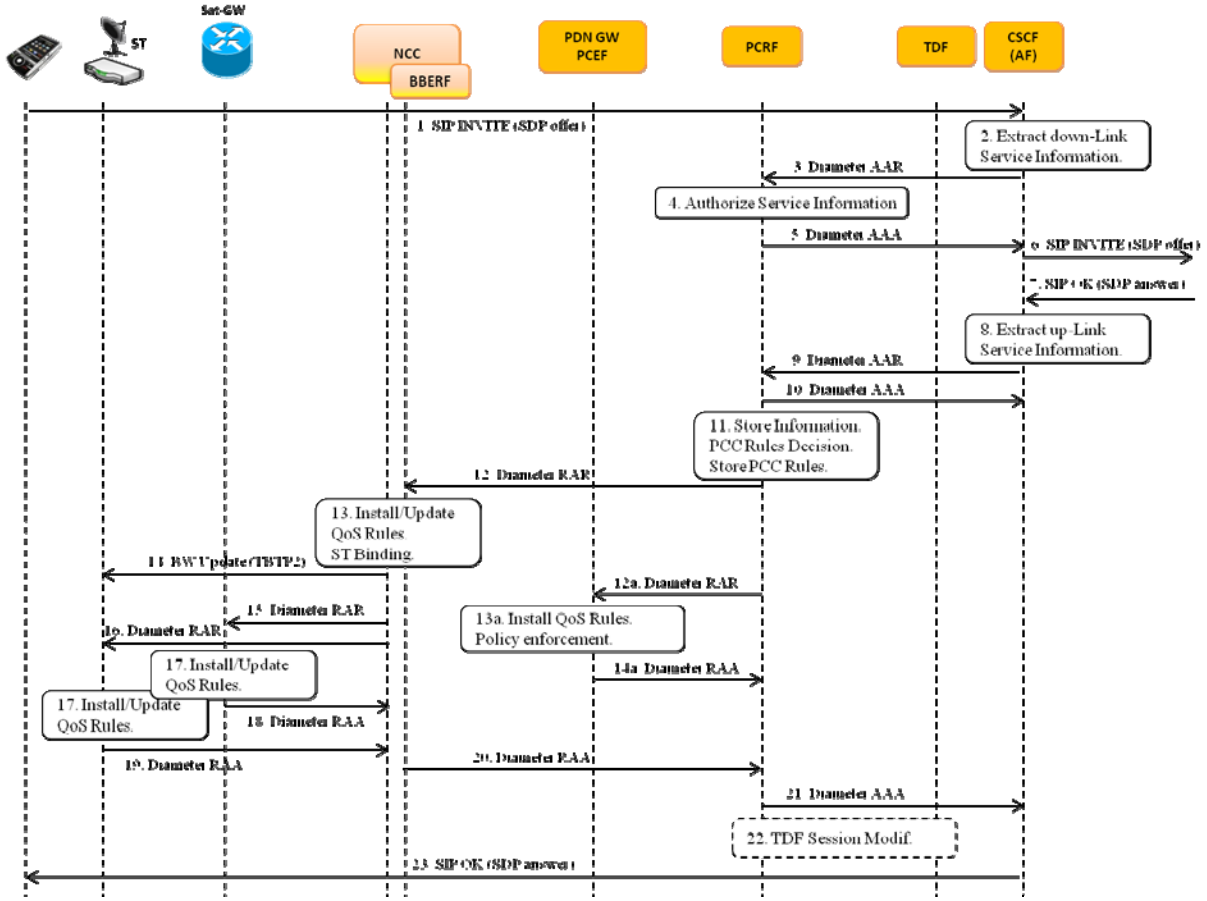


**Figure A.17: Detailed Signalling flow for IMS SIP call**

# Annex B: COMSEC recommendations

This clause presents the list of technical issues that occur from using Virtual Private Network (VPN) technologies in broadband satellite systems. The issues are described from the perspective of the satellite system integrator/operator point of view. The present document also proposes solutions to the technical issues, and provides recommendation and guidelines for efficient deployment of secure VPNs in broadband satellite systems.

Three control cases are defined:

- **Case 1:** The satellite system integrator/operator has control on both ends of the VPN or at least is able to recommend the VPN type, or the installation of features (e.g. Performance Enhancing Proxies, PEPs), or appropriate configurations on both sides of the VPN.
- **Case 2:** The satellite system integrator/operator has no control on both VPN sides. This case means that the satellite system integrator/operator cannot choose the VPN technology, or cannot configure or modify the VPN devices, or cannot install or recommend the installations of PEPs before VPN processing.
- **Case 3:** Case 3 is a mixture of Case 1 and Case 2. Here, the satellite system integrator/operator controls one end of the VPN but not the other. Hence, he is usually not able to recommend/chose the VPN technology but on one side of the VPN he is able to install or recommend the installation of (integrated) PEPs before VPN processing, or recommend configurations, etc.

## B.1 Issues with Performance Enhancing Proxies in secure VPNs

Figure B.1 illustrates the normal TCP operation as well as the interception performed by a TCP acceleration PEP. The deployment of transport layer PEPs is not an issue for TLS/SSL-based protection, which leaves the transport layer accessible.

In normal TCP operation, TCP data is acknowledge by the receiver after successful reception, meaning that it takes a round trip RTT1 until data is acknowledged. Since RTT1 is high in case a geostationary satellite link is involved, the bandwidth delay product limitation of TCP could be reached as mentioned previously.

A PEP is only able to perform TCP acceleration in case it can send a faked TCP ACK packet successful to the TCP sender well before the original TCP ACK, resulting in a round trip time of RTT2, which is lower than RTT1.

This has to be the function of the PEP independent whether the PEP splits the TCP session or just performs TCP ACK spoofing.

In the following, we will give reasons why this PEP function is not possible on IPsec-protected data:

- IPsec encryption: In case the TCP data is IPsec-encrypted, the PEP is unable to see the TCP header. Hence, it is not able to generate a TCP ACK message belonging to the respective TCP data. The PEP is even not able to determine the right TCP port and whether it is TCP data at all.
- IPsec without encryption: In case the TCP data is not encrypted but IPsec integrity protection is deployed, the PEP is able to see the TCP header. It is able to generate an appropriate TCP ACK message but the PEP is unable to perform IPsec integrity protection without knowing the IPsec key. When sending the fakes TCP ACK message towards the sender, the IPsec GW will drop it since, due to normal security policies, only IPsec protected data is allowed.

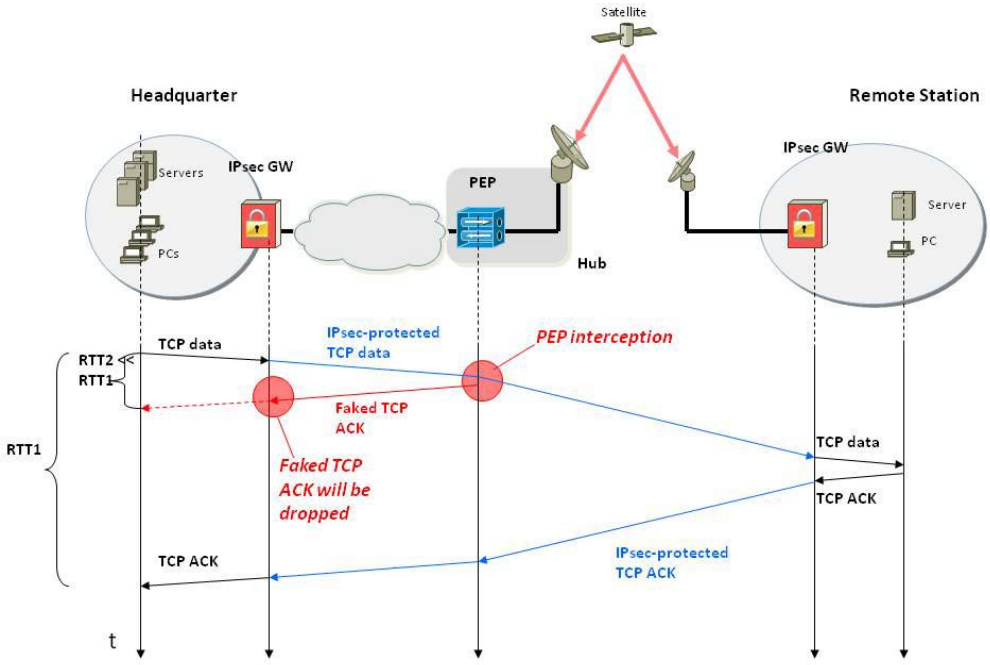


**Figure B.1: Manipulation of IPSec data**

## B.1.1 Possible solutions

### B.1.1.1 Positioning the distributed PEPs outside the VPN channel

Provided having control case 1and having IPsec in tunnel mode, a straight forward solutions is to place the PEP functions in the path not subject to VPN protection, e.g. before VPN processing at the sender side and after VPN processing at the receiver side. Deploying the PEP process outside the VPN channel allows the PEP functions to access the headers and payload data in scope to enhance performance. The network architecture for using TCP acceleration via a distributed PEP solution is illustrated in Figure B.2. The PEPs on both sides have full access to the TCP layer and are able to split the TCP connection to use an enhanced transport protocol over the satellite link. In some network architectures PEPs do payload compression before the data enters the VPN tunnel.

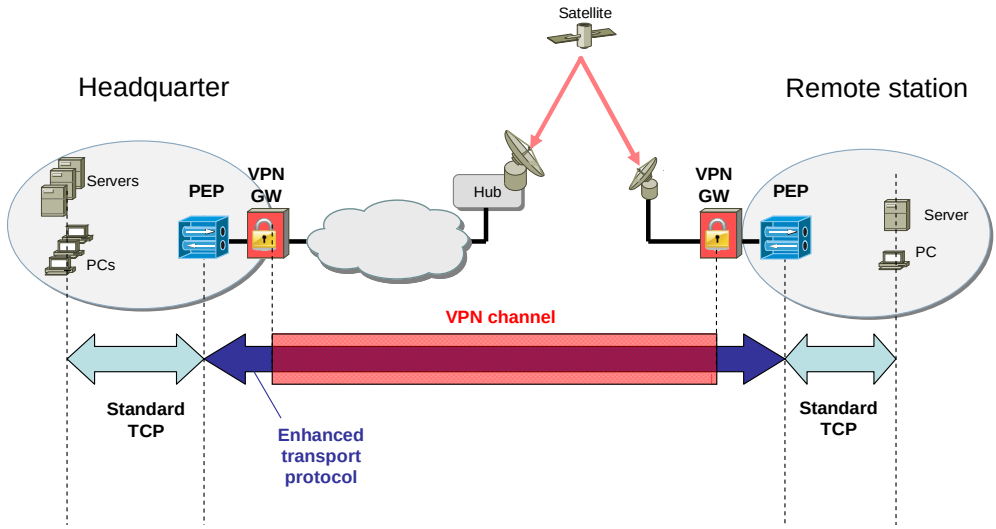


**Figure B.2: Distributed PEP positioned outside the VPN channel**

### B.1.1.2 Positioning the integrated PEP outside the VPN channel

Here we have to distinguish between an application layer PEP (e.g. HTTP cache or DNS cache) and a transport layer PEP, i.e. a PEP that is splitting TCP connections:

Application layer PEP:

Application layer PEPs like a web cache or a DNS cache are realized as integrated PEP. They are usually placed close to the hosts using the PEP to have short transmission times in case of a cache hit. Since the PEP usually does not know the VPN secret key, the only option is to place the PEP outside the VPN channel so that the respective protocol headers are accessible. The PEP terminates the application session between client and server and establishes a new session to the server. This network architecture is given in Figure B.3.

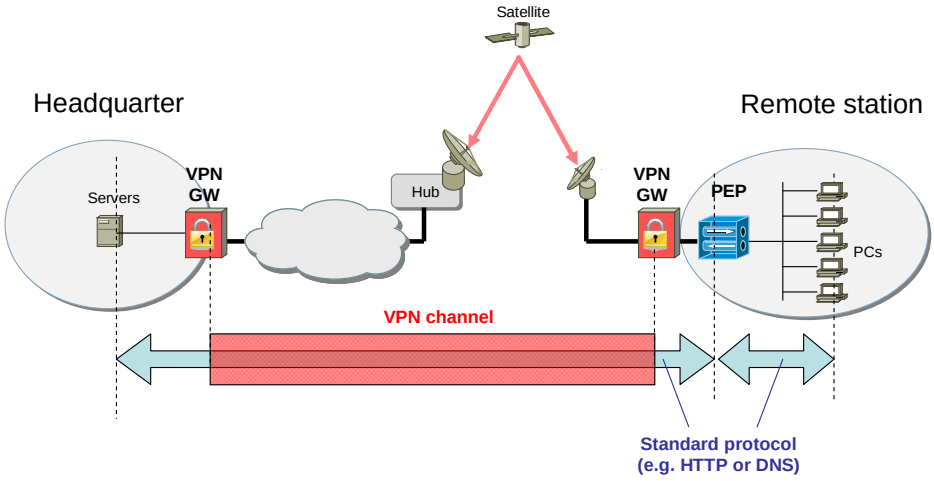


**Figure B.3: Integrated PEP positioned outside the VPN channel**

Integrated transport layer PEP:

An integrated transport layer PEP is usually based on a TCP splitting approach. It acts as the legal TCP receiver towards the TCP sender and terminates the TCP session by intercepting the TCP establishment. Afterwards, it establishes a new TCP session with the original TCP receiver. Hence, the TCP session is split into two parts. The usual deployment of an integrated PEP is at the satellite hub, as illustrated in Figure B.4. This has some advantages:

- In case the PEP is located at the satellite hub, the path between TCP sender and PEP is usually based on a high-speed terrestrial link so the data transfer towards the PEP can be high-speed without the issues of satellite links (high RTT, higher packet loss).
- Splitting the TCP session means that both TCP sessions have a lower RTT than the original TCP session.
- In case the second part between PEP and TCP client is a satellite connection, an enhanced TCP version can be used as long as this is compatible with the TCP version used at the TCP clients.

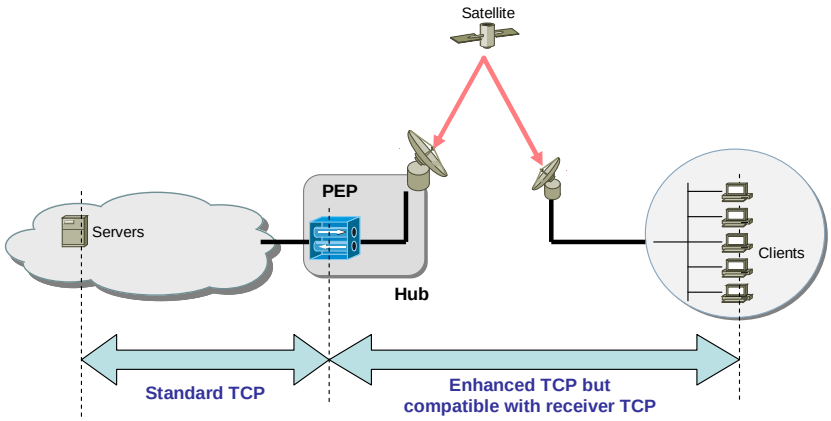


**Figure B.4: Integrated transport layer PEP at the hub**

Since operating at the transport layer, there are no issues with the deployment of VPNs based on TLS/SSL. However, in case of the deployment of a network layer VPN end-to-end the PEP is unable to access the transport layer without knowing the secret key, which is usually not acceptable for the user (from a security and management point of view). Hence, the only possibility is to position the PEP outside the VPN tunnel, which is just possible in case of IPsec is used in tunnel mode. Because the transport mode is used for end-to-end encryption, so there is no way to deploy the PEP outside the VPN tunnel. Two deployment options are possible:

- Positioning of the integrated PEP at the Headquarter: In case TCP downloads are performed from the headquarter (e.g. the headquarter of a company or organization) to the remote station, the PEP can be installed at the headquarter.
- Positioning of the integrated PEP at the Remote Station: In case TCP transfer is performed from the remote station to the headquarter (e.g. uploading of documents, videos, pictures, etc.), the PEP can be installed at the remote station.

### B.1.1.3 Deployment of SSL/TLS-aware proxies

There are some solutions available to accelerate applications even when the VPN channel is protected by SSL/TLS. However, these solutions require control case 1 or 3.

The deployment of SSL/TLS does not prevent TCP acceleration. However, some other performance improvements useful in satellite networks are not possible, e.g. caching and compression. In this clause, the deployment of caching in case of using HTTPS is discussed.

A HTTP caching proxy caches the web content (e.g. just certain media or complete web pages) when a webpage is requested the first time to have it available when the same media or webpage is requested once again, either by the same or a different user connected to the proxy. There are HTTP caching proxy implementations available that support HTTPS/TLS. Thereby, the HTTPS/TLS connection between web browser and web server is split at the proxy.

There are at least two preconditions given for this solution:

1) This function is not transparent for the user, i.e. the user has to explicitly configure the address of the caching proxy in its browser.
2) The user has to trust the proxy and the organization controlling it since in principle the proxy would be able to redirect requests to malicious servers. Hence, in most cases the proxy will be installed and controlled by the organization the user belongs to.

Taking these preconditions into account, this solution is just possible in case of having control case 1 and control case 3. Of course, this solution is the more beneficial the more often the same web content is requested: either a single user is requesting web content several times or several users are connected to the proxy and have a similar browsing behaviour and interest. In order to save the time required for transmitting cached content over the satellite link, the best place for the HTTPS cache is at the remote side, as illustrated in Figure B.5.

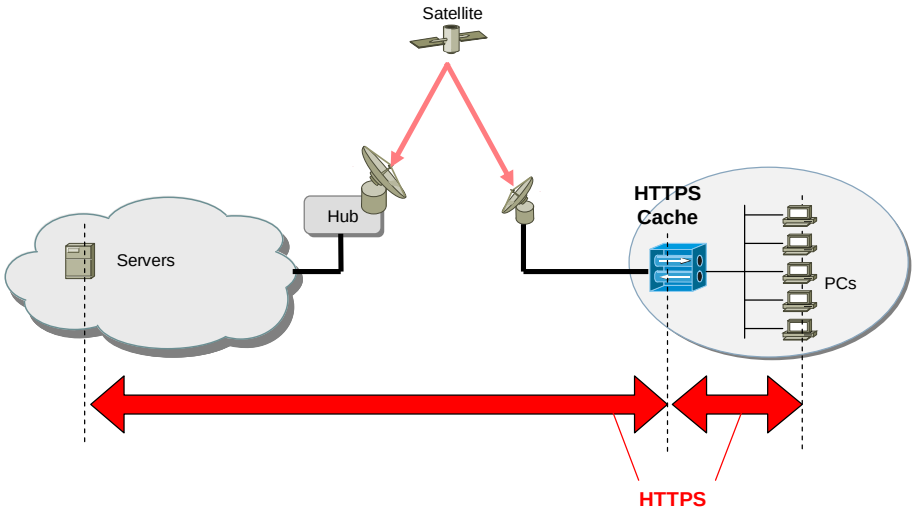


**Figure B.5: HTTPS caching proxy**

### B.1.1.4 Selection of transport layer or application layer VPN methods

Security mechanisms that protect the content above the transport layer (e.g. TLS/SSL or application layer security) allow PEPs on the protected path to perform TCP acceleration. Hence, in order to better support PEP deployment in satellite networks, a principle solution would be to choose VPN methods that operate above the transport layer instead of network layer VPNs.

However, there are various constraints that influence the choice of VPN method and the user is usually not free to select a VPN method that fits best. Therefore, it is usually not possible to switch from a network layer VPN solution to a transport layer or application layer VPN method just to allow PEP deployment. In the following, some reasons are given for keeping a network layer VPN solution:

- Higher level of security: A network layer VPN solution protects the fields of the transport layer and upper layer, e.g. port numbers. Furthermore, in case of IPsec AH, even some fields of the IP header are protected against manipulation. Changing to SSL/TLS means to weaken security.
- Missing security support in applications: A network layer VPN solution is usually deployed to have a single secure channel for all applications independent of the application. Switching to a SSL/TLS-based solutions or application layer security requires having security support in all applications of interest, which may not be given.
- Client/server model: TLS/SSL and also some application specific security features are based on client/server models, where a client starts the connection with one server. IPsec does not demand for a client/server model but is based on a peer to peer relationship. Depending on the scenario in scope, a client/server model may not be usable.

In summary, transport and higher layer security mechanisms are appropriate when possible to be deployed, but usually it is not possible to replace network layer VPNs by transport layer VPNs or application layer security.

## B.2 QoS enforcement issues in secure VPNs

With TLS and with IPsec in transport mode, QoS enforcement is not affected by VPN processing. As in IPsec tunnel mode the packet's original 5-tuple flow identifier is now replaced by the one of the VPN GW, QoS enforcement of IPsec-protected packets using those fields is not possible. This is because the original IP header (including the DSCP field) is replaced by the IP header generated at the VPN GW.

Figure B.6 illustrates the problem of Quality of Service (QoS) enforcement for tunnel mode IPsec-protected traffic.

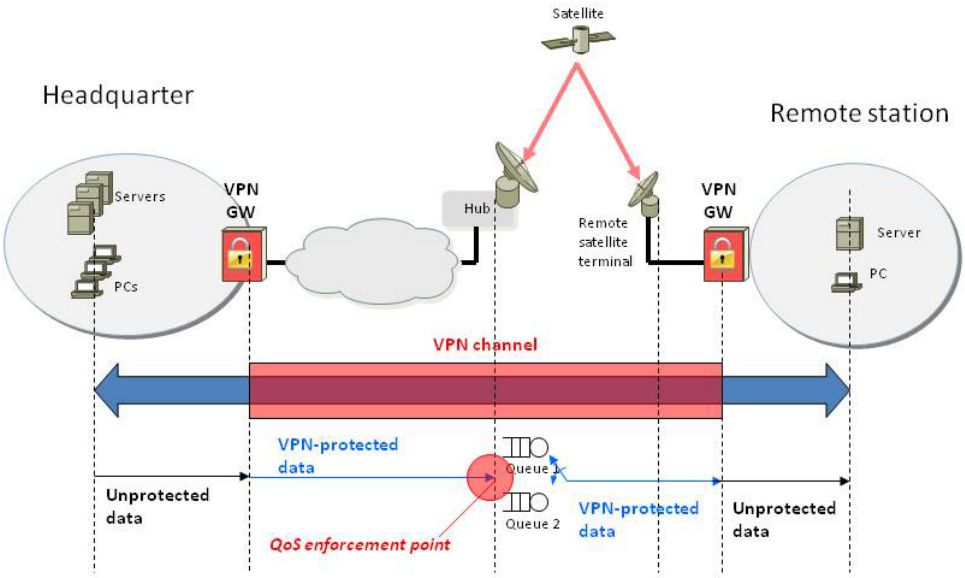


**Figure B.6: QoS enforcement issue with IPsec in tunnel mode**

## B.2.1 Possible solutions

### B.2.1.1 Copying DSCP field from inner to outer header

In the construction of the outer IP header, [i.53] specifies that the contents of DS field in the inner header should be copied to the outer header of a tunnel mode IPsec packet. It is applicable for both IPv4 and IPv6. Figure B.7 illustrates this process for ESP in tunnel mode.

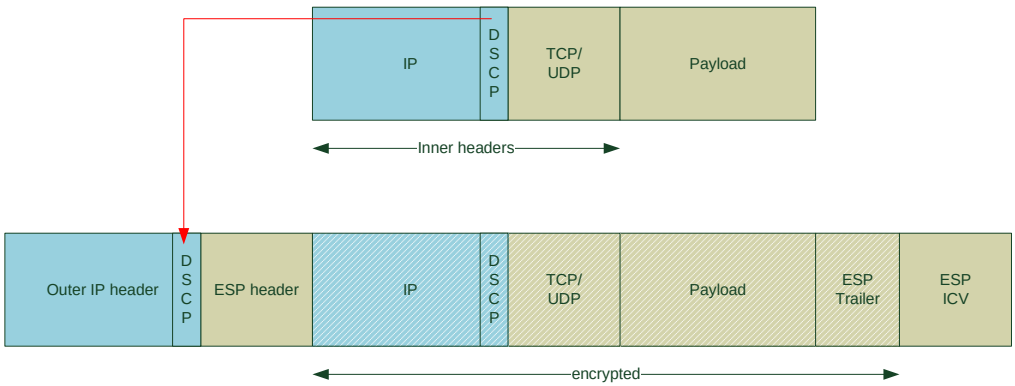


**Figure B.7: Copying of DSCP value from the inner to the outer header**

In order to classify packets based on the DSCP field, the inner DSCP field should be marked as close to the traffic source as possible, such that the correct value is visible to the VPN GW for further mapping to the outer header. This could either be directly at the end node/application or at the VPN GW before the packet is protected by IPsec. For example, the user VoIP application sets the DSCP value that corresponds to DiffServ Assured Forwarding (AF) traffic class (see [i.30]). Alternatively, the VPN GW could also set this field before it applies IPsec processing to the packet.

As the outer DSCP value reflects the original (the one of the inner header), it in turn reflects the intended QoS treatment of the packet. The QoS enforcement point at the hub or at the remote satellite terminal can then use the DSCP value to classify packets and deliver them into the proper transmission queues.

Some remarks are worth noting:

- Although this solution is mandated by IPsec, it might be undesirable in certain scenarios due to the security requirements. Copying the DSCP to the outer header means disclosing some information on the traffic flow characteristic, and thus potentially enabling a malicious party to perform traffic-analysis-based attacks.

- Because the solution has to be implemented at the user end-points (either at the user terminal or at the VPN GW), its implementation requires a control case 1.
- The mapping between the PHB / QoS policy and the DSCP value has to be agreed between the satellite operator and the end user. This can be achieved either by using the standard-defined DSCP values and PHB, or through a dedicated Service Level Agreement (SLA) established between the satellite operator and the user.
- For packets that have to pass through the Internet before arriving at the satellite operator's QoS enforcement point, it is important to ensure that none of the intermediate router modifies the DS field such that it would cause a different treatment of the packet in the satellite link.

# Annex C:
# Impact of random access on TCP behaviour

## C.1 TCP delay variation and packet misordering

A system that switches traffic flows from RA to DA channel could result in a change of the delay, or introduce variation of delay. This is especially the case for size-based network queuing. These considerations are most important for A-DAMA Top-Up and Back-Up use of RA with DAMA, since it is in these cases that the traffic may be divided between multiple physical layer transmission queues.

Sudden changes in delay could adversely impact the TCP RTT measurements, potentially resulting in expiry of the RTO and hence an unwanted congestion response. This is not expected to be a significant effect when using a modern TCP implementation. Where the change in delay is not accompanied by loss, the effect of spurious retransmissions may be reduced using methods such as the Eifel algorithms [i.54], [i.55] in the TCP sender or Forward RTO-recovery [i.56].

The impact of delay variation depends on the application. Most TCP applications, such as web browsing, are tolerant to small (<RTT) delay variations [i.57]. On the other hand, performance of real time applications such as VoIP can be affected by delay variation [i.58], resulting in loss at the input to a speech codec, or adversely impacting the round-trip estimator used to scale the playout buffer.

The switch of traffic (or unbalanced loading) from a RA to DA channel could result in reordering of packets at TCP receiver when the return path implements a queuing method that does not preserve per-flow order. Such methods could be motivated by the desire to use the RA channel for short packets, where it has the best possibility of reducing the queuing delay.

Examples of such methods include size-based queuing algorithms, such as ACKs-first scheduling and variants of the shortest-packet first algorithm. These methods have been used in early routers to reduce the queuing delay that can result when ACKs are queued behind larger data segments sent over low capacity links. The methods can produce unusual pathologies when used with TCP, and their impact depends on the traffic pattern, as described in the examples below. In these examples size-based network queuing is considered for a path that comprises a capacity-constrained return link and high-speed forward link using standard queuing:

- When the return link carries only ACKs, the forward path does not benefit from these methods.
- When the return link predominantly carries ACKs with occasional Data segments and ACKs, the forward link will observe decreased delay. Bi-directional flows, but may experience slower cwnd growth (misordered ACKs in Data segments do not inflate cwnd in modern TCP). Overall there may be benefit for forward transfers, which decreases as the volume of Data increases. The return link data performance may not be appreciably impacted, since ACKs are generally much smaller than data.
- When the return link predominantly carries Data segments with occasional ACKs, the forward link will benefit. But Shortest-First queuing can result in reordering of bursts when there are variable-sized Data segments, which can trigger Fast Retransmission and reduce performance.
- A return link that carries only Data, the forward path does not benefit from these methods, and the return path is not significantly impacted by ACK-First queuing. The impact of Shortest-First queuing depends on the viability of Data segments, but is not recommended, since it can result in unpredictable behavior with specific applications (e.g. block-oriented data transfers that typically send full-sized segments, but periodically send small segments at the end of each block).

The use of size-based queuing was common in early packet networks, but is not recommended for use in the general Internet, since it can lead to erratic application performance.

Reordering can also significantly impact the opening of the cwnd at the sender, which is important to the performance of short-lived flows. Excessive reordering beyond the SuPACKThreshold (currently 3 segments) will trigger fast retransmission and fast recovery, with resulting impact on the cwnd and ssthresh, and should generally be avoided in networks supporting TCP.

The effects of reordering may be mitigated by adapting the queuing algorithm to avoid simultaneous use of the RA and DA channels by a packet flow.

**Table C.1: Impact of integrated RA-DAMA at higher layers**

| MAC Layer | Network (IP) Layer | Transport (TCP) Layer | Application layer |
|---|---|---|---|
| ***RA Channel*** | ***Benefit at Higher Layers*** | | |
| No access delay | Lower Round Trip Delay (RTD) – Fast network response | Faster delivery & acknowledgement of application data | Better QoS performance for short interactive applications |
| **RA-DAMA** | **Impact/Issues at Higher Layers** | | |
| 1. Random packet losses will occur on RA channel | 1a. Lost packets are not recovered at IP layer | 1a. TCP RTO mechanisms are triggered. Delay depends on initial RTO value, prompt & accurate estimation of RTT. | 1a. QoS depends on how fast TCP can deliver the connection & data requests. |
| | 1b. Random packet losses cannot be differentiated from congestion losses | 1b. A spurious congestion signal is triggered, affecting TCP sender initial cwnd & ssthresh values. The impact depends on how conservatively TCP responds to congestion. | 1b. TCP congestion control affects QoS as small values of initial cwnd & ssthresh increase response time due to more round trips. |
| 2. Maximum RA bitrate is low (due to high cost) compared to DAMA | 2a. Packet reordering occurs if short packets are sent on RA & large packets on DAMA | 2a. TCP prematurely triggers fast retransmit/fast recovery if serious reordering occurs. TCP mechanisms are available to detect spurious retransmissions. | 2a. Packets are reordered by TCP. Delay & jitter components are experienced by application. |
| | 2b. Variable packet delay if large packets are switched/transmitted on RA and DAMA | 2b. TCP RTT estimation may be inaccurate leading to premature RTO. The impact may be insignificant if RTT is much longer than transmission time. | 2b. Some applications are sensitive to jitter (e.g. VoIP) but short interactive applications are more tolerant. |

# C.2 Responsiveness of standard TCP

The core principle for TCP congestion control is that loss of packets is regarded as a potential source of congestion. When packet loss is detected, TCP therefore activates its congestion control algorithms, as defined in [i.59].

Operating TCP over RA can result in loss of control or request packets at the beginning of a transmission. This may trigger overly conservative behavior, even though there is no congestion. In this case, response time is mostly affected by the state of RTO and IW variables, which are dynamically set according to TCP conservative principles.

## C.2.1 Reduced initial RTO

It is recommended in [i.60] reducing the initial RTO of TCP from a previous value of 3 seconds to 1 second, unless the SYN or SYN-ACK is lost, in which case the default RTO is reverted to 3 seconds before data transmission begins. The lower RTO value was found to be sufficient for more than 97,5 % of connections, while implication of spurious retransmissions for few connections with RTT longer than 1 second is modest. More significantly, the new value is small enough to ensure timely recovery from packet losses occurring before an RTT sample is taken. Hence this new standard enhances TCP response time in case of initial packet loss.

However, RTO loss recovery activates congestion control thus causing the TCP sender to be overly conservative during non-congestion periods. In particular, the following two state variables are affected:

- Congestion Window (cwnd) - This is set to Loss Window. Slow start restarts with cwnd of only 1 segment.
- Slow-start Threshold (ssthresh) - The ssthresh is set to around 2 segments i.e. max (FlightSize/2, 2*SMSS).

## C.2.2 Early loss recovery

Modern TCP uses algorithms to detect and recover from loss within the shortest possible time, usually before an RTO has expired. These mechanisms include:

### C.2.2.1 Fast Retransmit and Fast Recovery

The Fast Retransmit/Fast Recovery algorithm allows a TCP receiver to send an immediate duplicate ACK when it receives an out-of-order segment that confirms data is held waiting for a particular byte number. The TCP sender uses the Fast Retransmit algorithm [i.59] to detect and repair loss, based on incoming ACKs. The arrival of 3 duplicate ACKs acts as an indication that a segment has been lost. Hence fast recovery of lost segment can be performed without incurring an RTO. SACK is a widely employed enhancement to this method – but has little impact on the first few packets of a flow.

### C.2.2.2 Limited transmit

When the flight size is less than 4 segments, Fast Retransmit cannot be used, because there will never be more sufficient Dup ACKs to trigger the method. The Limited Transmit algorithm [i.61] allows an additional outstanding segment to be sent upon receiving each Dup ACK (increasing flight size). This eventually triggers Fast Retransmit, when 3 Dup ACKs may be induced after a loss.

### C.2.2.3 Early retransmit

The Limited Transmit algorithm cannot trigger Fast Retransmit if TCP sender does not have additional outstanding segments to send up to required amount (e.g. a burst is limited to 3 segments or less, such many web page requests). To solve this problem, the Early Retransmit algorithm [i.62] calculates a new value (ER_thresh) that determines number of DUP ACKs needed to trigger Fast Retransmit based on outstanding unsent data. In the case of a burst of 3 segments, this method reduces the number of DUP ACKs required to trigger Fast Retransmit to only 2.

## C.3.3 Redundant TCP SYNs

The TCP standard specifies sending a single initial SYN packet and waiting for an ACK. The SYN is only retransmitted after the RTO period, when a loss is assumed thus delaying actual start of data transmission.

SYN duplication is a proposed technique that could improve TCP responsiveness when the initial SYN packet is lost. One way to achieve this is by setting initial RTO smaller than the actual path RTT. It has been argued that since general-purpose networks are designed for large traffic flows, it is reasonably safe to be aggressive when sending short flows [i.63]. The RTO retransmit timer can be set low e.g. 100ms, or even less, if packets are not too close together to share the same fate. (In doing so, it is important to verify that the implementation does not reset *ssthresh* when performing a SYN retransmission.)

In a satellite system, duplication of the initial SYN could save time at the expense of using an additional transmission burst (e.g. in RA channel). However the overhead of SYN duplication may not be a significant, because the additional SYN packet is only 40B (without compression).

In general, TCP responsiveness is affected primarily by how fast a client is able to deliver the connection and data requests. Thus, additional redundancy should be employed for duplicated SYN or request packets rather than subsequent confirmations.

Delayed ACK or other proposed ACK Congestion Control mechanisms [i.64] may offset the possible extra load due on the RA channel.

There are however concerns with resending a SYN, since some clients use this as method for detecting whether the end host supports a particular function. For example, only the initial SYN may carry some TCP options, and loss of this SYN could significantly change the operation of the remaining connection. One notable example is dual-stack systems, where an IPv6 sender may revert to IPv4 for the second SYN, since it is assumed that the server failed to respond to the initial network layer request. Care should be taken to avoid such effects impacting the user.

## C.3.4 Changing TCP RTT/RTO estimation

[i.65] specifies a method that does not sample the RTT during the three-way handshake (3WHS) when using a large IW, because delay changes can result once a session is established. One example is when there is a significant time to serialize a data packet on a narrowband link, where seeding the RTO based on an RTT of a small SYN or SYN-ACK packets would likely underestimate the RTT for larger data packets.

A proposal from Google [i.66] recommends sampling the RTT during the 3WHS and seeding the RTO regardless of the size of the IW. The main reasoning for this proposed reversal of practice is the prominence of faster links in the Internet suffering noticeable latency while waiting for an RTO compared to the benefits of a shorter RTT. Seeding the RTO with correct RTT sampled after SYN and SYN-ACK exchange has been suggested to improve TCP responsiveness in the case of loosing subsequent packets during the handshake. However, it is important to note that this proposal was made before [i.60] became the standard for computing TCP retransmission timer.

Reduction of the initial RTO from 3 seconds to 1 second may reduce the urgency of this particular proposal. The proposal could still be useful for links with very short RTT, but could raise issues on slow links or links that rely on DA or RA methods. Further research is required to judge the applicability to the general Internet.

### C.3.5 Sending data with TCP SYN

TCP Fast Open (TFO) [i.67] is another proposal that would allow data to be carried in the SYN or SYN-ACK packets and consumed by the receiving end during the initial connection handshake. This provides a saving of up to one full RTT compared to standard TCP, requiring a 3WHS to complete before data can be exchanged. Data on SYN behavior was allowed in [i.68] but TFO would additionally allow data to be delivered to the application before the 3WHS has completed.

In the proposed method, the server side uses a security cookie to authenticate a client initiating a TFO connection thus addressing previous data integrity concerns caused by dubious SYN packets. This avoids the pitfalls of earlier methods, such as T/TCP. However, it requires an additional exchange between client and server at the beginning of a connection for requesting the fast open cookie, which should also be expired by the server after some time. TFO is somewhat limited, as it is more applicable for applications that have temporal locality on client and server connections.

There are concerns with sending data on SYN such as a client choosing IP version (IPv6 or IPv4) that is not supported at the server, or starting with an unknown size of the Maximum Segment Size (MSS) for the link. Additionally, there is no sequence number protection hence the packet is more vulnerable to attacks. For the moment this remains a topic of research, if accepted this proposal would significantly increase the size of a TCP SYN, which may impact usage of the RA channel.

### C.3.6 Increasing TCP Initial Window

An increase of IW from 1 to 3 segments has been widely deployed, motivated by the desire to improve Fast Retransmit.

A recent proposal from Google argues for increasing IW further to at least ten segments (about 15KB) for speedy completion of short TCP transfers in one RTT. Furthermore, reduction in total transfer time for data greater than 4KB up to 4 RTTs is possible. Preliminary experiments by Google show benefits in reducing object transfer times at moderate cost in terms of increased congestion and associated packet losses. This analysis did not explore the potential collateral damage on other flows that share a bottleneck where the large IW is continuously used.

Google has also recommended that TCP implementations refrain from resetting IW to one segment unless there have been multiple SYN or SYN-ACK retransmissions, or true loss detection has been made [i.62]. The current standard [i.65] specifies resetting IW to 1 on loosing even a single control packet. However, considering [i.60] reduction of initial RTO from 3 seconds to 1 second, it is possible to unnecessarily penalize connections with high RTT values (e.g. satellite links).

A key argument to be assessed is that there is little or no experience of using a larger IW on other flows that share a constrained path. The likely impact on real-time flows (voice, video) may be significant if many flows use a larger IW. This is therefore an area of current research, and a topic where standards are expected within the TCPM working group of the IETF.

**Table C.2: Recent proposals to enhance responsiveness of standard TCP**

| TCP Mechanism | Standard RFC | Proposed Enhancement |
|---|---|---|
| Retransmission Timeout (RTO) | RTO not seeded during three-way handshake [i.60] | Google (Seeding RTO with RTT sampled during three-way handshake) |
| Initial Congestion Window (IW) | Maximum initial window of 3 segments [i.65] | Google (Increasing TCP IW from 3 to 10 segments) |
| Loss Window (LW) | Reduce IW to 1 segment on loss of packet during three-way handshake [i.65] | Google (Refrain from resetting IW to LW upon loss of packet during the three-way handshake) |
| Initial SYN control packet | TCP sender sends one initial SYN to start connection [i.68] | Damon Wischik (Setting initial RTO smaller than RTT e.g. to duplicate SYN) |
| ACK control | Delayed ACK [i.59] | [i.64] (ECN-marked ACK packets) |
| Data on SYN | [i.68] forbids the receiver to deliver the data to the application until 3WHS is completed. | Google TCP Fast Open (allows the receiver to deliver the data to the application during 3WHS) |

# History

| Document history | | |
|---|---|---|
| V1.1.1 | April 2014 | Publication |
| | | |
| | | |
| | | |
| | | |