

International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.9904

(10/2012)

SERIES G: TRANSMISSION SYSTEMS AND MEDIA,
DIGITAL SYSTEMS AND NETWORKS

Access networks – In premises networks

**Narrowband orthogonal frequency division
multiplexing power line communication
transceivers for PRIME networks**

Recommendation ITU-T G.9904



ITU-T G-SERIES RECOMMENDATIONS

TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

INTERNATIONAL TELEPHONE CONNECTIONS AND CIRCUITS	G.100–G.199
GENERAL CHARACTERISTICS COMMON TO ALL ANALOGUE CARRIER-TRANSMISSION SYSTEMS	G.200–G.299
INDIVIDUAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON METALLIC LINES	G.300–G.399
GENERAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON RADIO-RELAY OR SATELLITE LINKS AND INTERCONNECTION WITH METALLIC LINES	G.400–G.449
COORDINATION OF RADIOTELEPHONY AND LINE TELEPHONY	G.450–G.499
TRANSMISSION MEDIA AND OPTICAL SYSTEMS CHARACTERISTICS	G.600–G.699
DIGITAL TERMINAL EQUIPMENTS	G.700–G.799
DIGITAL NETWORKS	G.800–G.899
DIGITAL SECTIONS AND DIGITAL LINE SYSTEM	G.900–G.999
MULTIMEDIA QUALITY OF SERVICE AND PERFORMANCE – GENERIC AND USER-RELATED ASPECTS	G.1000–G.1999
TRANSMISSION MEDIA CHARACTERISTICS	G.6000–G.6999
DATA OVER TRANSPORT – GENERIC ASPECTS	G.7000–G.7999
PACKET OVER TRANSPORT ASPECTS	G.8000–G.8999
ACCESS NETWORKS	G.9000–G.9999
In premises networks	G.9900–G.9999

For further details, please refer to the list of ITU-T Recommendations.

Recommendation ITU-T G.9904

Narrowband orthogonal frequency division multiplexing power line communication transceivers for PRIME networks

Summary

Recommendation ITU-T G.9904 contains the physical layer (PHY) and data link layer (DLL) specification for PRIME narrowband orthogonal frequency division multiplexing (OFDM) power line communication transceivers for communications via alternating current and direct current electric power lines over frequencies in the CENELEC A band.

This Recommendation uses material from Recommendations ITU-T G.9955, ITU-T G.9956 and ITU-T G.9956 Amendment 1; specifically using material from Annex B of ITU-T G.9955, Annex B of ITU-T G.9956 and ITU-T G.9956 Amendment 1. New technical material has not been introduced in this version.

The control parameters that determine spectral content, power spectral density (PSD) mask requirements and the set of tools to support the reduction of the transmit PSD can be found in Recommendation ITU-T G.9901.

History

Edition	Recommendation	Approval	Study Group
1.0	ITU-T G.9904	2012-10-29	15

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure, e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2013

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

Table of Contents

	Page
1	Scope 1
2	References..... 1
3	Definitions 2
4	Abbreviations and acronyms 3
5	Conventions 6
6	General description 6
6.1	General description of the architecture..... 6
7	Physical layer..... 7
7.1	Introduction 7
7.2	Overview 7
7.3	PHY parameters..... 8
7.4	Preamble, header and payload structure 8
7.5	Convolutional encoder..... 11
7.6	Scrambler..... 12
7.7	Interleaver..... 12
7.8	Modulation 13
7.9	Electrical specification of the transmitter..... 15
7.10	PHY service specification 16
8	Data link layer specifications..... 30
8.1	Overview 30
8.2	Addressing..... 31
8.3	MAC functional description 34
8.4	MAC PDU format 54
8.5	MAC service access point 77
8.6	MAC procedures 95
8.7	Automatic repeat request (ARQ)..... 110
9	Convergence layer 114
9.1	Overview 114
9.2	Common part convergence sublayer (CPCS)..... 114
9.3	NULL specific service convergence sublayer (NULL SSCS) 116
9.4	IPv4 specific service convergence sublayer (IPv4 SSCS) 117
9.5	IEC 61334-4-32 specific service convergence sublayer (IEC 61334-4-32 SSCS) 129
9.6	IPv6 service-specific convergence sublayer (IPv6 SSCS) 133

	Page
10 Management plane.....	148
10.1 Introduction	148
10.2 Node management	149
10.3 Firmware upgrade.....	162
10.4 Management interface description	179
10.5 List of mandatory PIB attributes	184
Annex A – EVM and SNR calculation	187
Annex B – MAC layer constants	188
Annex C – Convergence layer constants	189
Appendix I – Examples of CRC	190
Appendix II – Interleaving matrices	191

Recommendation ITU-T G.9904

Narrowband orthogonal frequency division multiplexing power line communication transceivers for PRIME networks

1 Scope

This Recommendation contains the physical layer (PHY) and data link layer (DLL) specification for PRIME narrowband orthogonal frequency division multiplexing (OFDM) power line communication transceivers for communications via alternating current and direct current electric power lines over frequencies below 500 kHz. This Recommendation supports indoor and outdoor communications over low-voltage lines, medium-voltage lines, through transformer low-voltage to medium-voltage and through transformer medium-voltage to low-voltage power lines in both urban and in long distance rural communications. This Recommendation addresses grid to utility meter applications, advanced metering infrastructure (AMI), and other 'Smart Grid' applications such as the charging of electric vehicles, home automation and home area networking (HAN) communications scenarios.

This Recommendation removes the control parameters that determine spectral content, power spectral density (PSD) mask requirements, and the set of tools to support reduction of the transmit PSD, all of which have been moved to [ITU-T G.9901].

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

- [ITU-T G.9901] Recommendation ITU-T G.9901 (2012), *Narrowband orthogonal frequency division multiplexing power line communication transceivers – Power spectral density specification*.
- [IEC 61334-4-1] IEC 61334-4-1 Ed.1996, *Distribution automation using distribution line carrier systems – Part 4: Data communication protocols – Section 1: Reference model of the communication system*.
- [IEC 61334-4-32] IEC 61334-4-32 Ed.1996, *Distribution automation using distribution line carrier systems – Part 4: Data communication protocols – Section 32: Data link layer – Logical link control (LLC)*.
- [IEC 61334-4-511] IEC 61334-4-511 Ed. 2000, *Distribution automation using distribution line carrier systems – Part 4-511: Data communication protocols – Systems management – CIASE protocol*.
- [IEC 61334-4-512] IEC 61334-4-512, Ed. 1.0:2001, *Distribution automation using distribution line carrier systems – Part 4-512: Data communication protocols – System management using profile 61334-5-1 – Management Information Base (MIB)*.
- [IEEE 802-2001] IEEE Std 802-2001 (R2007), *IEEE Standard for Local and Metropolitan Area Networks. Overview and Architecture*.
- [IETF RFC 768] IETF RFC 768 (1980), *User Datagram Protocol (UDP)*.

- [IETF RFC 791] IETF RFC 791 (1981), *Internet Protocol, DARPA Internet Program, Protocol Specification*.
- [IETF RFC 793] IETF RFC 793 (1981), *Transmission Control Protocol (TCP)*.
- [IETF RFC 1144] IETF RFC 1144 (1990), *Compressing TCP/IP Headers for Low-Speed Serial Links*.
- [IETF RFC 2131] IETF RFC 2131 (1997), *Dynamic Host Configuration Protocol (DHCP)*.
- [IETF RFC 2460] IETF RFC 2460 (1998), *Internet Protocol, Version 6 (IPv6) Specification*.
- [IETF RFC 2462] IETF RFC 2462 (1998), *IPv6 Stateless Address Autoconfiguration*.
- [IETF RFC 2464] IETF RFC 2464 (1998), *Transmission of IPv6 Packets over Ethernet Networks*.
- [IETF RFC 3022] IETF RFC 3022 (2001), *Traditional IP Network Address Translator (Traditional NAT)*.
- [IETF RFC 3315] IETF RFC 3315 (2003), *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*.
- [IETF RFC 4291] IETF RFC 4291 (2006), *IP Version 6 Addressing Architecture*.
- [IETF RFC 4862] IETF RFC 4862 (2007), *Ipv6 Stateless Address Autoconfiguration*.
- [IETF RFC 6282] IETF RFC 6282 (2011), *Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks*.
- [EN 50065-1] CENELEC EN 50065-1 (2011), *Signalling on low-voltage electrical installations in the frequency range 3 kHz to 148,5 kHz – Part 1: General requirements, frequency bands and electromagnetic disturbances*.
- [PUB 197] NIST FIPS PUB 197 (2001), *Advanced Encryption Standard (AES)*.
- [SP 800-38A] NIST SP 800-38A (2001), *Recommendation for Block Cipher Modes of Operation. Methods and Techniques*.
- [SP 800-57] NIST SP 800-57 (2007), *Recommendation for Key Management – Part 1: General (Revised)*.

3 Definitions

3.1 Terms defined elsewhere

None.

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

- 3.2.1 base node:** The master node which controls and manages the resources of a subnetwork.
- 3.2.2 beacon slot:** The location of the beacon PDU within a frame.
- 3.2.3 destination node:** A node that receives a frame.
- 3.2.4 downlink:** Data travelling in the direction from the base node towards the service nodes.
- 3.2.5 level (PHY layer):** When used in the physical layer (PHY) context, it implies the transmit power level.
- 3.2.6 level (MAC layer):** When used in the medium access control (MAC) context, it implies the position of the reference device in switching hierarchy.

3.2.7 MAC frame: A composite unit of abstraction of time for channel usage. A MAC frame is comprised of one or more beacons, one SCP, and zero or one CFP. The transmission of the beacon by the base node acts as a delimiter for the MAC frame.

3.2.8 neighbour node: Node A is neighbour node of node B if A can directly transmit to and receive from B.

3.2.9 node: Any one element of a subnetwork which is able to transmit to and receive from other subnetwork elements.

3.2.10 PHY frame: The set of OFDM symbols and preamble which constitute a single PHY layer protocol data unit (PPDU).

3.2.11 preamble: The initial part of a PHY frame, used for synchronization purposes.

3.2.12 registration: The process by which a service node is accepted as member of a subnetwork and allocated an LNID.

3.2.13 service node: Any one node of a subnetwork which is not a base node.

3.2.14 source node: A node that sends a frame.

3.2.15 subnetwork: A set of elements that can communicate by complying with this Recommendation and share a single base node.

3.2.16 subnetwork address: The property that universally identifies a subnetwork. It is its base node EUI-48 address.

3.2.17 switching: Providing connectivity between nodes that are not neighbour nodes.

3.2.18 unregistration: The process by which a service node leaves a subnetwork.

3.2.19 uplink: Data travelling in the direction from the service node towards the base node.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

AC	Alternating Current
ADC	Analogue-to-Digital Converter
AES	Advanced Encryption Standard
AMM	Advanced Meter Management
ARQ	Automatic Repeat request
ATM	Asynchronous Transfer Mode
BER	Bit Error Rate
BPDU	Beacon PDU
BPSK	Binary Phase Shift Keying
BSI	Beacon Slot Information
CFP	Contention-Free Period
CID	Connection Identifier
CL	Convergence Layer
CIMTUSize	Convergence layer Maximum Transmit Unit Size
CPCS	Common Part Convergence Sublayer

CRC	Cyclic Redundancy Check
CSMA-CA	Carrier Sense Multiple Access-Collision Avoidance
D8PSK	Differential Eight-Phase Shift Keying
DBPSK	Differential Binary Phase Shift Keying
DHCP	Dynamic Host Configuration Protocol
DPSK	Differential Phase Shift Keying (general)
DQPSK	Differential Quaternary Phase Shift Keying
DSK	Device Secret Key
ECB	Electronic Code Book
EMA	Exponential Moving Average
ENOB	Effective Number Of Bits
EUI-48	48-bit Extended Unique Identifier
EVM	Error Vector Magnitude
FCS	Frame Check Sequence
FEC	Forward Error Correction
FFT	Fast Fourier Transform
FU	Firmware Upgrade
GK	Generation Key
GPDU	Generic MAC PDU
HCS	Header Check Sum
IFFT	Inverse Fast Fourier Transform
IGMP	Internet Group Management Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
KDIV	Key Diversifier
LCID	Local Connection Identifier
LFSR	Linear Feedback Shift Register
LLC	Logical Link Control
LNID	Local Node Identifier
LSID	Local Switch Identifier
LWK	Local Working Key
MAC	Medium Access Control
MK	Master Key
MLME	MAC Layer Management Entity
MOL	Maximal Output Level
MPDU	MAC Protocol Data Unit
MSB	Most Significant Bit

MSDU	MAC Service Data Unit
MSPS	Million Samples Per Second
MTU	Maximum Transmission Unit
NAT	Network Address Translation
NID	Node Identifier
NSK	Network Secret Key
OFDM	Orthogonal Frequency Division Multiplexing
PDU	Protocol Data Unit
PHY	Physical Layer
PIB	PLC Information Base
PLC	Power Line Communications
PLME	PHY Layer Management Entity
PNPDU	Promotion Needed PDU
PPDU	PHY Protocol Data Unit
ppm	Parts per million
PRBS	Pseudo-Random Binary Sequence
PSD	Power Spectral Density
PSDU	PHY Service Data Unit
QoS	Quality of Service
SAP	Service Access Point
SAR	Segmentation And Reassembly
SCP	Shared-Contention Period
SCRC	Secure CRC
SDU	Service Data Unit
SEC	Security
SID	Switch Identifier
SNA	Subnetwork Address
SNK	Subnetwork Key (corresponds to either REG.SNK or SEC.SNK)
SNR	Signal to Noise Ratio
SP	Security Profile
SSCS	Service Specific Convergence Sublayer
SWK	Subnetwork Working Key
TCP	Transmission Control Protocol
TOS	Type Of Service
UI	Unique Identifier
USK	Unique Secret Key
VJ	Van Jacobson

5 Conventions

Binary numbers are indicated by the prefix '0b' followed by the binary digits, e.g., '0b0101'. Hexadecimal numbers are indicated by the prefix '0x', e.g., '0x0F'.

Mandatory requirements are indicated with 'shall' in the main body of this Recommendation.

Optional requirements are clearly indicated. If an option is incorporated in an implementation, it shall be applied as specified in this Recommendation.

roof (.) denotes rounding to the closest higher or equal integer.

floor (.) denotes rounding to the closest lower or equal integer.

$A \bmod B$ denotes the remainder (from 0, 1, ..., B-1) obtained when an integer A is divided by an integer B.

6 General description

This Recommendation provides a solution for power line communications in the bands specified by the main body and Annex C of [ITU-T G.9901] using OFDM modulation. This solution focuses on providing a very robust communication channel in applications such as advanced meter management (AMM). The target transmission rate is in the order of tens of kilobits per second.

6.1 General description of the architecture

Figure 6-1 depicts the communication layers and the scope of this Recommendation.

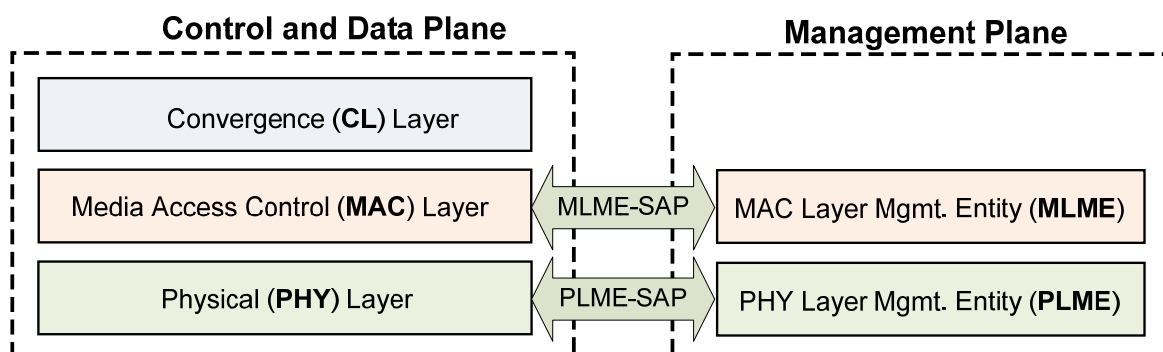


Figure 6-1 – Reference model of ITU-T G.9904 protocol layers

The convergence layer (CL) classifies traffic associating it with its proper MAC connection. This layer performs the mapping of any kind of traffic to be properly included in MAC service data units (MSDUs). It may also include compression functions. Several SSCs are defined to accommodate different kinds of traffic into MSDUs.

The MAC layer provides core MAC functionalities of system access, bandwidth allocation, connection establishment/maintenance and topology resolution.

The PHY layer transmits and receives MPDUs between neighbour nodes using orthogonal frequency division multiplexing (OFDM). OFDM is chosen as the modulation technique mainly because of the following:

- its inherent adaptability in the presence of frequency selective channels (which are quite common but unpredictable, due to narrowband interference or unintentional jamming);
- its robustness to impulsive noise, resulting from extended symbol duration and use of FEC;

- its capacity for achieving high spectral efficiencies with simple transceiver implementations.

7 Physical layer

7.1 Introduction

This clause specifies the physical layer (PHY) entity for an OFDM based PLC communications scheme in the bands specified in the main body and Annex C of [ITU-T G.9901].

Differential modulation is used with three possible constellations: DBPSK, DQPSK or D8PSK. Thus, theoretical uncoded speeds of around 47 kbit/s, 94 kbit/s, and 141 kbit/s (without a cyclic prefix overhead) could be obtained.

An additive scrambler is used to avoid the occurrence of long sequences of identical bits.

Finally, rate $\frac{1}{2}$ convolutional coding will be used along with bit interleaving. This can be disabled by higher layers if the channel is good enough and higher throughputs are needed.

7.2 Overview

Figure 7-1 shows the PHY layer transmitter block diagram.

On the transmitter side, the PHY layer receives an MPDU from the MAC layer and generates a PHY frame. If decided by higher layers, the PPDU after the CRC block is convolutionally encoded and then interleaved (however, it will always be scrambled). The output is differentially modulated using a DBPSK, DQPSK or D8PSK scheme. The next step is OFDM, which comprises the inverse fast Fourier transform (IFFT) block and the cyclic prefix generator.




Figure 7-1 – PHY layer transmitter

The structure of the PHY frame is shown in Figure 7-2. Each PHY frame starts with a preamble lasting 2.048 ms, followed by a number of OFDM symbols, each lasting 2.24 ms. The first two OFDM symbols carry the PHY frame header. The PHY header is also generated as described in clause 7.4. The remaining M OFDM symbols carry payload, generated as described in clause 7.4. The value of M is signalled in the PHY header, and is at most equal to 63.




Figure 7-2 – PHY frame format

7.3 PHY parameters

PHY parameters are defined in clause C.2 of [ITU-T G.9901].

There are parameters which depend on the modulation of each OFDM subcarrier.

Table 7-1 shows the PHY data rate during payload transmission, and the maximum MSDU length for various modulation and coding combinations.

Table 7-1 – PHY data rates for each modulation with the FEC on and off

	DBPSK		DQPSK		D8PSK	
	On	Off	On	Off	On	Off
Convolutional code (½)	On	Off	On	Off	On	Off
Information bits per subcarrier, N_{BPSC}	0.5	1	1	2	1.5	3
Information bits per OFDM symbol, N_{BPS}	48	96	96	192	144	288
Raw data rate (kbit/s approx)	21.4	42.9	42.9	85.7	64.3	128.6
Maximum MSDU length (in bits) with 63 symbols	3016	6048	6040	12096	9064	18144
Maximum MSDU length (in bytes) with 63 symbols	377	756	755	1512	1133	2268

Table 7-2 shows the modulation and coding scheme and the size of the header portion of the PHY frame

Table 7-2 – Header parameters

	DBPSK
Convolutional code (½)	On
Information bits per subcarrier, N_{BPSC}	0.5
Information bits per OFDM symbol, N_{BPS}	42

It is strongly recommended that all frequencies used to generate the OFDM transmit signal come from one single frequency reference. The system clock shall have a maximum tolerance of ± 50 ppm, including ageing.

7.4 Preamble, header and payload structure

7.4.1 Preamble

The preamble is used at the beginning of every PPDU for synchronization purposes. In order to provide a maximum of energy, a constant envelope signal is used instead of OFDM symbols. There is also a need for the preamble to have frequency agility that will allow synchronization in the presence of frequency selective attenuation and of course, excellent aperiodic autocorrelation properties are mandatory. A linear chirp signal meets all the above requirements. The waveform of the preamble is defined as:

$$S_{CH}(t) = A \cdot \text{rect}(t/T) \cdot \cos\left[2\pi\left(f_0 t + 1/2\mu t^2\right)\right]$$

where T , f_0 , f_s , and μ are defined in clause C.3 of [ITU-T G.9901].

$\text{rect}(\cdot)$ function is defined as:

$$\begin{aligned} \text{rect}(t) &= 1, & 0 < t < 1 \\ \text{rect}(t) &= 0, & \text{otherwise} \end{aligned}$$

7.4.2 Pilot structure

The two OFDM symbols comprising the PHY header shall contain 13 pilot subcarriers which can be used to estimate the sampling start error and the sampling frequency offset.

For subsequent OFDM symbols, one pilot subcarrier is used to provide a phase reference for frequency domain DPSK demodulation.

Pilot subcarrier frequency allocation is shown in Figures 7-3 and 7-4, where P_i is the i -th pilot subcarrier and D_j is the j -th data subcarrier.





Figure 7-3 – Pilot and data subcarrier allocation (OFDM symbols vs subcarriers)



G.9955(11)_FB-5

Figure 7-4 – Pilot and data subcarrier frequency allocation inside the header

Pilot subcarriers shall be BPSK modulated by a pseudo-random binary sequence (PRBS) to prevent the generation of spectral lines. The phase of the pilot subcarriers is controlled by the PRBS sequence which is a cyclic extension of the 127 bit sequence given by:

$\text{Pref}_{0..126} =$

{0,0,0,0,1,1,1,0,1,1,1,1,0,0,1,0,1,1,0,0,1,0,0,1,0,0,0,0,0,0,1,0,0,0,1,0,0,1,1,0,0,0,1,0,1,1,1,0,1,0,1,1,0,1,1,0,0,0,0,1,1,0,0,1,1,0,1,0,1,1,0,1,1,0,0,0,0,1,1,0,0,1,1,0,1,0,1,0,1,0,0,1,1,1,0,0,1,1,1,1,0,1,1,0,1,0,0,0,0,1,0,1,0,1,0,1,1,1,1,1,0,1,0,0,1,0,1,0,0,0,1,1,0,1,1,1,0,0,0,1,1,1,1,1,1,1,1,1,1,1}

where '1' means 180° phase shift and '0' means 0° phase shift. One bit of the sequence will be used per pilot subcarrier, starting with the first pilot subcarrier in the first OFDM symbol, then the next pilot subcarrier, and so on. The same process is used for the second OFDM symbol. For subsequent OFDM symbols, one element of the sequence is used for the pilot subcarrier (see Figure 7-3).

The PRBS sequence can be generated by the scrambler defined in Figure 7-5 when the "all-ones" initial state is used.




Figure 7-5 – LFSR for use in pilot sequence generation


The loading of the PRBS sequence shall be initiated at the start of every PPDU, just after the preamble.

7.4.3 PHY header and payload

The PHY header is composed of two OFDM symbols which are always sent using DBPSK modulation and FEC (convolutional coding) 'On'. However the payload is DBPSK, DQPSK or D8PSK modulated, depending on the configuration by the MAC layer. The MAC layer selects the modulation scheme, e.g., using information from errors in previous transmissions to the same receiver(s), or by using the SNR feedback. Thus, the system will then configure itself dynamically to provide the best compromise between throughput and efficiency in the communication. This includes deciding whether or not FEC (convolutional coding) is used.

The first two OFDM symbols in the PPDU corresponding to the PHY header are composed of 84 data subcarriers and 13 pilot subcarriers. After the PHY header, each OFDM symbol in the payload carries 96 data subcarriers and one pilot subcarrier. Each data subcarrier carries 1, 2 or 3 bits.

The bit stream from each field shall be sent MSB first. See Figure 7-6.



G.9955(11)_FB-7

Figure 7-6 – PPDU: PHY header and payload (bits transmitted before encoding)

PHY HEADER: Each PPDU contains both PHY and MAC header information. It is composed of the following fields:

- **PROTOCOL:** contains the transmission scheme of the payload. Added by the PHY layer.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
DBPSK	DQPSK	D8PSK	RES	DBPSK_F	DQPSK_F	D8PSK_F	RES	RES	RES	RES	RES	RES	RES	RES	RES

Figure 7-7 – "PROTOCOL" field of the PPDU

- Where RES means "reserved" and the suffix "_F" means FEC is 'On'.
- **LEN:** defines the length of the payload (after coding) in OFDM symbols. Added by the PHY layer.
- **PAD_LEN:** defines the length of the PAD field (before coding) in bytes. Added by the PHY layer.
- **MAC_H:** MAC layer header. It is included inside the header symbols to protect the information contained. The MAC header is generated by the MAC layer and only the first 54 bits of the MAC header are embedded in the PHY header.
- **CRC_Ctrl:** the $CRC_Ctrl(m)$, $m = 0..7$, contains the CRC checksum over PROTOCOL, LEN, PAD_LEN and MAC_H field (PD_Ctrl). The polynomial form of PD_Ctrl is expressed as follows:

$$\sum_{m=0}^{69} PD_{Ctrl}(m)x^m$$

The checksum is calculated as follows: the remainder of the division of PD_Ctrl by the polynomial $x^8 + x^2 + x + 1$ forms $CRC_Ctrl(m)$, where $CRC_Ctrl(0)$ is the LSB. The generator polynomial is the well-known CRC-8-ATM. Some examples are shown in Appendix I.


- **FLUSHING_H:** flushing bits needed for convolutional decoding. All bits in this field are set to zero to reset the convolutional encoder. Added by the PHY layer.

PAYLOAD:

- **MSDU:** uncoded MAC layer service data unit.
- **FLUSHING_P:** flushing bits needed for convolutional decoding. All bits in this field are set to zero to reset the convolutional encoder. This field only exists when FEC is 'On'.
- **PAD:** In order to ensure that the number of (coded) bits generated in the payload fills an integer number of OFDM symbols, pad bits shall be added to the payload before encoding. All pad bits shall be set to zero.

7.5 Convolutional encoder

The uncoded PHY stream can be convolutionally encoded to form the encoded PHY stream. The encoder is a rate $\frac{1}{2}$ convolutional encoder with constraint length $K = 7$ and code generator "polynomials" 1111001 and 1011011. At the beginning, the encoder state is set to zero. At the end of the transmission of either the header or the payload, zeroes shall be inserted to flush the encoder (8 zeros for the PHY header and 6 zeros for the payload). The bit generated by the first code generator is first output. The block diagram of the encoder is shown in Figure 7-8.



G.9955(11)_FB-8

Figure 7-8 – Convolutional encoder

7.6 Scrambler


The scrambler block randomizes the bit stream so it reduces the crest factor at the output of the IFFT. Scrambling shall always be performed.

The scrambler block performs an XOR of the input bit stream using a pseudo noise sequence pn obtained by a cyclic extension of the 127 element sequence given by:

$\text{Pref}_{0..126} =$

{0,0,0,0,1,1,1,0,1,1,1,1,0,0,1,0,1,1,0,0,1,0,0,1,0,0,0,1,0,0,0,0,0,1,0,0,0,1,0,0,1,1,0,0,0,1,0,1,1,1,0,1,0,1,1,0,1,1,0,1,1,0,0,0,0,1,1,0,0,1,1,0,1,0,1,0,1,0,1,0,1,1,0,0,0,0,1,1,0,0,1,1,0,1,0,1,0,1,1,1,0,0,1,1,1,0,1,1,0,1,1,0,0,0,0,1,0,1,0,1,0,1,1,1,1,0,1,0,0,1,0,1,0,0,0,1,1,0,1,1,1,0,0,0,1,1,1,1,1,1,1,1,1,1}

The PRBS sequence can be generated by the scrambler defined in Figure 7-9 when the "all-ones" initial state is used.



G.9955(11)_FB-9

Figure 7-9 – LFSR for use in the scrambler block

Loading of the sequence pn shall be initiated at the start of every PPDU, just after the preamble.

7.7 Interleaver

Because of the frequency fading (narrowband interference) of typical power line channels, OFDM subcarriers are generally received at different amplitudes. Deep fades in the spectrum may cause groups of subcarriers to be less reliable than others, thereby causing bit errors to occur in bursts rather than be randomly scattered. If (and only if) coding is used as described in clause 7.4, interleaving is applied to randomize the occurrence of bit errors prior to decoding. At the

transmitter, the coded bits are permuted in a certain way, which ensures that adjacent bits are separated by several bits after interleaving.

Let $N_{CBPS} = 2 \times N_{BPS}$ be the number of coded bits per OFDM symbol in the cases where convolutional coding is used. All coded bits shall be interleaved by a block interleaver with a block size corresponding to N_{CBPS} . The interleaver ensures that adjacent coded bits are mapped onto non-adjacent data subcarriers. Let $v(k)$, with $k = 0, 1, \dots, N_{CBPS}-1$, be the coded bits vector at the interleaver input. $v(k)$ is transformed into an interleaved vector $w(i)$, with $i = 0, 1, \dots, N_{CBPS}-1$, by the block interleaver as follows:

$$w((N_{CBPS}/s) \times (k \bmod s) + \text{floor}(k/s)) = v(k) \quad k = 0, 1, \dots, N_{CBPS}-1.$$

The value of s is determined by the number of coded bits per subcarrier, $N_{CBPSC} = 2 \times N_{BPSC}$. N_{CBPSC} is related to N_{CBPS} so that $N_{CBPS} = 96 \times N_{CBPSC}$ (payload) and $N_{CBPS} = 84 \times N_{CBPSC}$ (header)

$$s = 8 \times (1 + \text{floor}(N_{CBPSC}/2)) \text{ for the payload and}$$

$$s = 7 \text{ for the header.}$$

The de-interleaver performs the inverse operation. Hence, if $w'(i)$, with $i = 0, 1, \dots, N_{CBPS}-1$, is the de-interleaver vector input, the vector $w'(i)$ is transformed into a de-interleaved vector $v'(k)$, with $k = 0, 1, \dots, N_{CBPS}-1$, by the block de-interleaver as follows:

$$v'(s \times i - (N_{CBPS}-1) \times \text{floor}(s \times i / N_{CBPS})) = w'(i) \quad i = 0, 1, \dots, N_{CBPS}-1.$$

Descriptive tables showing index permutations can be found in Appendix II for reference.

7.8 Modulation

The PPDU payload is modulated as a multicarrier differential phase shift keying signal with one pilot subcarrier and 96 data subcarriers that comprise 96, 192 or 288 bits per symbol. The header is modulated DBPSK with 13 pilot subcarriers and 84 data subcarriers that comprise 84 bits per symbol.

The bit stream coming from the interleaver is divided into groups of M bits where the first bit of the group of M is the most significant bit (MSB).

Frequency domain differential modulation is performed. Figure 7-10 shows the DBPSK, DQPSK and D8PSK mapping.

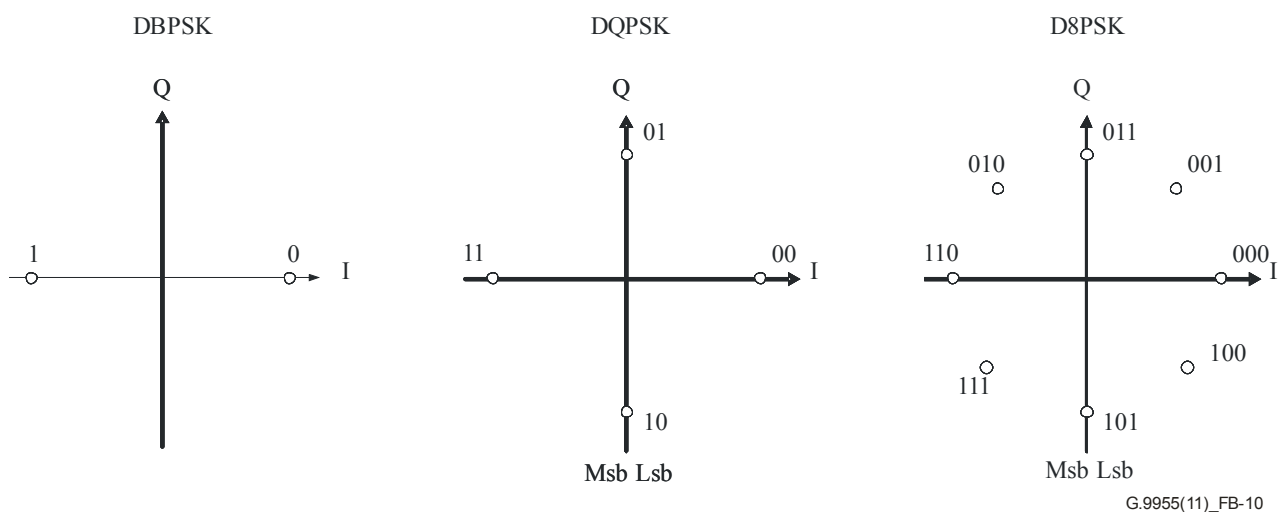


Figure 7-10 – DBPSK, DQPSK and D8PSK mapping

The next equation defines the M -ary DPSK constellation of M phases:

$$s_k = Ae^{j\theta_k}$$

where:

k is the frequency index representing the k -th subcarrier in an OFDM symbol.
 $k = 1$ corresponds to the phase reference pilot subcarrier.

s_k is the modulator output (a complex number) for the k -th given subcarrier.

θ_k stands for the absolute phase of the modulated signal obtained as follows:

$$\theta_k = (\theta_{k-1} + (2\pi/M)\Delta b_k) \bmod 2\pi$$

This equation applies for $k > 1$ in the payload, the $k = 1$ subcarrier being the phase reference pilot. When the header is transmitted, the pilot allocated in the k -th subcarrier is used as a phase reference for the data allocated in the $k+1$ -th subcarrier.

- $\Delta b_k \in \{0, 1, \dots, M-1\}$ represents the information coded in the phase increment, as supplied by the constellation encoder.
- $M = 2, 4, \text{ or } 8$ in the case of DBPSK, DQPSK or D8PSK, respectively.
 A represents the ring radius from the centre of the constellation.

The OFDM symbol can be expressed in mathematical form:

$$c_i(n) = \left\{ \sum_{k=86}^{182} s(k-85, i) \exp\left(\frac{j2\pi 2}{512}(n - N_{CP})\right) + \sum_{k=330}^{426} s(427-k, i) \exp\left(\frac{j2\pi 2}{512}(n - N_{CP})\right) \right\}$$

i is the time index representing the i -th OFDM symbol; $i = 0, 1, \dots, M+1$

n is the sample index; $48 \leq n \leq 559$ $s(k, i)$ is the complex value from the subcarrier modulation block and the symbol $*$ denotes complex conjugate.

If a complex 512-point IFFT is used, the 96 subcarriers shall be mapped as shown in Figure 7-11. The symbol '*' represents complex conjugate.

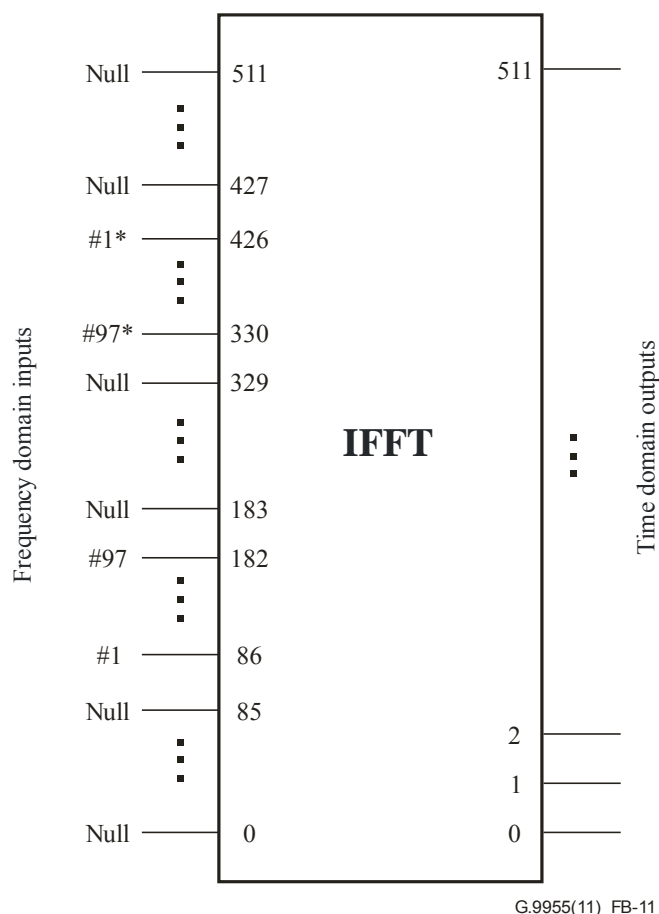


Figure 7-11 – Subcarrier mapping

After the inverse Fourier transform, the symbol is cyclically extended by 48 samples to create the cyclic prefix (N_{CP}).

7.9 Electrical specification of the transmitter

7.9.1 General

The following requirements establish the minimum technical transmitter requirements for interoperability and adequate transmitter performance. These requirements are specified in [ITU-T G.9901].

7.9.2 Transmit PSD

See clause C.4.2 of [ITU-T G.9901].

7.9.3 Error vector magnitude (EVM)

The quality of the injected signal with regard to the artificial mains network impedance shall be measured in order to validate the transmitter device. Accordingly, a vector analyser that provides EVM measurements (EVM meter) shall be used, see Annex A for an EVM definition. The test set-up described in Figures 4 and 6 of [EN 50065-1] shall be used in the case of single-phase devices and three-phase devices transmitting simultaneously on all phases, respectively.




Figure 7-12 – EVM meter (block diagram)

The EVM meter shall include a band pass filter with an attenuation of 40 dB at 50 Hz that ensures anti-aliasing for the analogue-to-digital converter (ADC).

The minimum performance of the ADC is 1 MSPS, 14-bit effective number of bits (ENOB). The ripple and the group delay of the band pass filter must be accounted for in EVM calculations.

7.9.4 Conducted disturbance limits

See clause C.4.3 of [ITU-T G.9901].

7.10 PHY service specification

7.10.1 General

The PHY shall have a single 20-bit free-running clock incremented in steps of 10 µsec. The clock counts from 0 to 1048575, then overflows back to 0. As a result the period of this clock is 10.48576 seconds. The clock is never stopped or restarted. Time measured by this clock is the one to be used in some PHY primitives to indicate a specific instant in time.

7.10.2 PHY data plane primitives

7.10.2.1 General

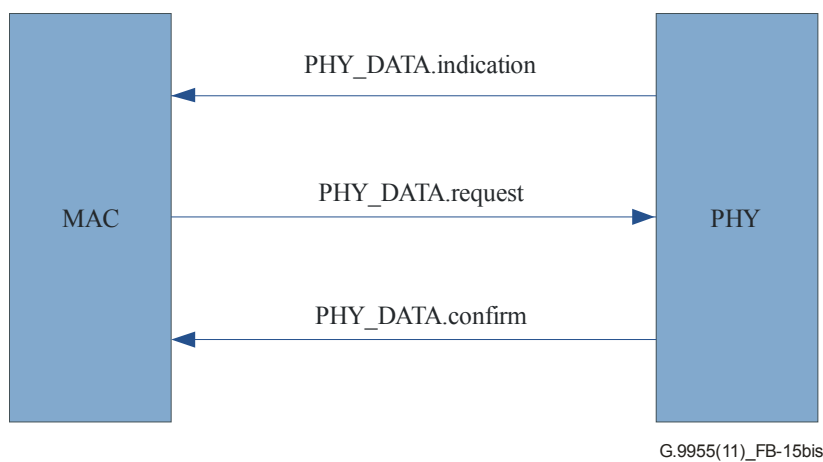


Figure 7-13 – Overview of PHY primitives

The request primitive is passed from the MAC to the PHY to request the initiation of a service.

The indication and confirm primitives are passed from the PHY to the MAC to indicate an internal PHY event that is significant to MAC. This event may be logically related to a remote service request or may be caused by an event internal to PHY.

7.10.2.2 PHY_DATA.request

7.10.2.2.1 Function

The PHY_DATA.request primitive is passed to the PHY layer entity to request the sending of a PPDU to one or more remote PHY entities using the PHY transmission procedures. It also allows setting the time at which the transmission shall be started.

7.10.2.2.2 Structure

The semantics of this primitive are as follows:

PHY_DATA.request{*MPDU*, *Length*, *Level*, *Scheme*, *Time*}.

The *MPDU* parameter specifies the MAC protocol data unit to be transmitted by the PHY layer entity. It is mandatory for implementations to byte-align the MPDU across the PHY-SAP. This implies 2 extra bits (due to the non-byte-aligned nature of the MAC layer header) to be located at the beginning of the header.

The *Length* parameter specifies the length of the MPDU in bytes. The length shall be 2 bytes long.

The *Level* parameter specifies the output signal level according to which the PHY layer transmits the MPDU. It may take one of eight values:

- 0: maximal output level (MOL)
- 1: MOL –3 dB
- 2: MOL –6 dB
- ...
- 7: MOL –21 dB.

The *Scheme* parameter specifies the transmission scheme to be used for the MPDU. It may have any of the following values:

- 0: DBPSK
- 1: DQPSK
- 2: D8PSK
- 3: not used
- 4: DBPSK + convolutional code
- 5: DQPSK + convolutional code
- 6: D8PSK + convolutional code
- 7: not used.

The *Time* parameter specifies the instant in time in which the MPDU has to be transmitted. It is expressed in tens of μsec and may take values from 0 to $2^{20} - 1$.

The *Time* parameter shall be calculated by the MAC, taking into account the current PHY time obtained by PHY_timer.get primitive. The MAC shall account for the fact that no part of the PPDU can be transmitted during beacon slots and CFP periods granted to other devices in the network. If the *Time* parameter is set such that these rules are violated, the PHY will return a fail in PHY_Data.confirm.

7.10.2.2.3 Use

The primitive is generated by the MAC layer entity whenever data is to be transmitted to a peer MAC entity or entities.

The reception of this primitive will cause the PHY entity to perform all the PHY-specific actions and pass the properly formed PPDU for transfer to the peer PHY layer entity or entities. The next transmission shall start when *Time* = Timer.

7.10.2.3 PHY_DATA.confirm

7.10.2.3.1 Function

The PHY_DATA.confirm primitive has only local significance and provides an appropriate response to a PHY_DATA.request primitive. The PHY_DATA.confirm primitive tells the MAC layer entity whether or not the MPDU of the previous PHY_DATA.request has been successfully transmitted.

7.10.2.3.2 Structure

The semantics of this primitive are as follows:

PHY_DATA.confirm{*Result*}.

The *Result* parameter is used to pass status information back to the local requesting entity. It is used to indicate the success or failure of the previous associated PHY_DATA.request. Some results will be standard for all implementations:

- 0: success
- 1: too late (the time for transmission has passed)
- 2: invalid length
- 3: invalid scheme
- 4: invalid level
- 5: buffer overrun
- 6: busy channel
- 7-255: proprietary

7.10.2.3.3 Use

The primitive is generated in response to a PHY_DATA.request.

It is assumed that the MAC layer has sufficient information to associate the confirm primitive with the corresponding request primitive.

7.10.2.4 PHY_DATA.indication

7.10.2.4.1 Function

This primitive defines the transfer of data from the PHY layer entity to the MAC layer entity.

7.10.2.4.2 Structure

The semantics of this primitive are as follows:

PHY_DATA.indication{*PSDU*, *Length*, *Level*, *Scheme*, *Time*}.

The *PSDU* parameter specifies the PHY service data unit as received by the local PHY layer entity. It is mandatory for implementations to byte-align the MPDU across the PHY-SAP. This implies two extra bits (due to the non-byte-aligned nature of the MAC layer header) to be located at the beginning of the header.

The *Length* parameter specifies the length of the received PSDU in bytes. The length shall be 2 bytes long.

The *Level* parameter specifies the signal level on which the PHY layer received the PSDU. It may take one of sixteen values:

- 0: ≤ 70 dBuV
- 1: ≤ 72 dBuV
- 2: ≤ 74 dBuV
- ...
- 15: > 98 dBuV

The *Scheme* parameter specifies the scheme with which the PSDU is received. It may have any of the following values:

- 0: DBPSK
- 1: DQPSK
- 2: D8PSK
- 3: not used
- 4: DBPSK + convolutional code
- 5: DQPSK + convolutional code
- 6: D8PSK + convolutional code
- 7: not used

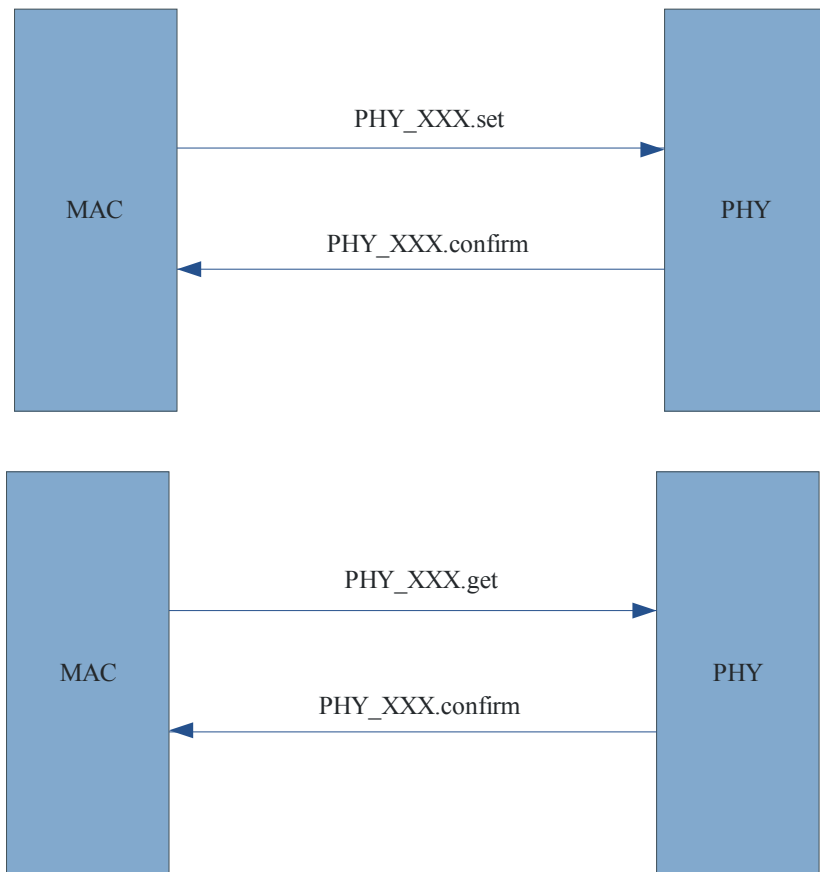
The *Time* parameter is the time of receipt of the preamble associated with the PSDU.

7.10.2.4.3 Use

The PHY_DATA.indication is passed from the PHY layer entity to the MAC layer entity to indicate the arrival of a valid PPDU.

7.10.3 PHY control plane primitives

Figure 7-14 shows the general structure of PHY control plane primitives. Each primitive may have "set", "get" or "confirm" fields. Table 7-3 lists the control plane primitives and the fields associated with each of them. Each row lists a control plane primitive. An "X" in a column indicates that the associated field is used in the primitive described in that row.



G.9955(11)_FB-15ter

Figure 7-14 – Overview of PHY control plane primitives

Table 7-3 – Fields associated with PHY control plane primitives

	set	get	confirm
PHY_AGC	X	X	X
PHY_Timer		X	X
PHY_CD		X	X
PHY_NL		X	X
PHY_SNR		X	X
PHY_ZCT		X	X

7.10.3.1 PHY_AGC.set

7.10.3.1.1 Function

The PHY_AGC.set primitive is passed to the PHY layer entity by the MAC layer entity to set the automatic gain mode of the PHY layer.

7.10.3.1.2 Structure

The semantics of this primitive are as follows:

PHY_AGC.set {*Mode*, *Gain*}.

The *Mode* parameter specifies whether or not the PHY layer operates in automatic gain mode. It may take one of two values:

- 0: Auto
- 1: Manual

The *Gain* parameter specifies the initial receiving gain in auto mode. It may take one of N values:

- 0: min_gain dB
- 1: min_gain + step dB
- 2: min_gain + 2 × step dB
- ...
- $N - 1$: min_gain + $(N - 1) \times$ step dB

where min_gain and N depend on the specific implementation. "step" is also an implementation issue but it shall not be more than 6 dB. The maximum Gain value min_gain + $(N - 1) \times$ step shall be at least 21 dB.

7.10.3.1.3 Use

The primitive is generated by the MAC layer when the receiving gain mode has to be changed.

7.10.3.2 PHY_AGC.get

7.10.3.2.1 Function

The PHY_AGC.get primitive is passed to the PHY layer entity by the MAC layer entity to get the automatic gain mode of the PHY layer.

7.10.3.2.2 Structure

The semantics of this primitive are as follows:

PHY_AGC.get{ }.

7.10.3.2.3 Use

The primitive is generated by the MAC layer when it needs to know the receiving gain mode that has been configured.

7.10.3.3 PHY_AGC.confirm

7.10.3.3.1 Function

The PHY_AGC.confirm primitive is passed to the MAC layer entity by the PHY layer entity in response to a PHY_AGC.set or PHY_AGC.get command.

7.10.3.3.2 Structure

The semantics of this primitive are as follows:

PHY_AGC.confirm {*Mode*, *Gain*}.

The *Mode* parameter specifies whether or not the PHY layer is configured to operate in automatic gain mode. It may take one of two values:

- 0: auto
- 1: manual

The *Gain* parameter specifies the current receiving gain. It may take one of N values:

- 0: min_gain dB
- 1: min_gain + step dB
- 2: min_gain + 2 × step dB
- ...
- $N - 1$: min_gain + $(N - 1) \times$ step dB

where min_gain and N depend on the specific implementation. The parameter step shall not be more than 6 dB. The maximum gain value min_gain + $(N - 1) \times$ step will be at least 21 dB.

7.10.3.4 PHY_Timer.get

7.10.3.4.1 Function

The PHY_Timer.get primitive is passed to the PHY layer entity by the MAC layer entity to get the time at which the transmission has to be started.

7.10.3.4.2 Structure

The semantics of this primitive are as follows:

PHY_Timer.get {}.

7.10.3.4.3 Use

The primitive is generated by the MAC layer to know the transmission start.

7.10.3.5 PHY_Timer.confirm

7.10.3.5.1 Function

The PHY_Timer.confirm primitive is passed to the MAC layer entity by the PHY layer entity in response to a PHY_Timer.get command.

7.10.3.5.2 Structure

The semantics of this primitive are as follows:

PHY_Timer.confirm {*Time*}

The *Time* parameter is specified in tens of microseconds. It may take values of between 0 and 220 – 1 .

7.10.3.6 PHY_CD.get

7.10.3.6.1 Function

The PHY_CD.get primitive is passed to the PHY layer entity by the MAC layer entity to look for the carrier detect signal. The carrier detection algorithm shall be based on preamble detection and header recognition (see clause 7.4).

7.10.3.6.2 Structure

The semantics of this primitive are as follows:

PHY_CD.get {}.

7.10.3.6.3 Use

The primitive is generated by the MAC layer when it needs to know whether or not the physical medium is free.

7.10.3.7 PHY_CD.confirm

7.10.3.7.1 Function

The PHY_CD.confirm primitive is passed to the MAC layer entity by the PHY layer entity in response to a PHY_CD.get command.

7.10.3.7.2 Structure

The semantics of this primitive are as follows:

PHY_CD.confirm {*cd*, *rss*i, *Time*, *header*}.

The *cd* parameter may take one of two values:

- 0: no carrier detected
- 1: carrier detected.

The *rss*i parameter is the received signal strength indication and refers to the preamble. It is only relevant when *cd* equals 1. It may take one of sixteen values:

- 0: ≤ 70 dBuV
- 1: ≤ 72 dBuV
- 2: ≤ 74 dBuV
- ...
- 15: > 98 dBuV

The *Time* parameter indicates the instant at which the present PPDU will finish. It is only relevant when *cd* equals 1. When *cd* equals 0, the time parameter will take a value of 0. If *cd* equals 1 but the duration of the whole PPDU is still not known (i.e., the header has not yet been processed), the header parameter will take a value of 1 and the time parameter will indicate the instant at which the header will finish, specified in tens of microseconds. In any other case the value of the time parameter is the instant at which the present PPDU will finish and it is specified in tens of microseconds. Time parameter refers to an absolute point in time so it is referred to the system clock.

The *header* parameter may take one of two values:

- 1: if a preamble has been detected but the duration of the whole PPDU is not yet known from decoding the header
- 0: in any other case.

7.10.3.8 PHY_NL.get

7.10.3.8.1 Function

The PHY_NL.get primitive is passed to the PHY layer entity by the MAC layer entity to get the floor noise level value.

7.10.3.8.2 Structure

The semantics of this primitive are as follows:

PHY_NL.get {}.

7.10.3.8.3 Use

The primitive is generated by the MAC layer when it needs to know the noise level present in the power line.

7.10.3.9 PHY_NL.confirm

7.10.3.9.1 Function

The PHY_NL.confirm primitive is passed to the MAC layer entity by the PHY layer entity in response to a PHY_NL.get command.

7.10.3.9.2 Structure

The semantics of this primitive are as follows:

PHY_NL.confirm {*noise*}.

The *noise* parameter may take one of sixteen values:

0: ≤ 50 dBuV

1: ≤ 53 dBuV

2: ≤ 56 dBuV

...

15: > 92 dBuV

7.10.3.10 PHY_SNR.get

7.10.3.10.1 Function

The PHY_SNR.get primitive is passed to the PHY layer entity by the MAC layer entity to get the value of the signal-to-noise ratio, defined as the ratio of the measured received signal level to noise level of the last received PPDU. The calculation of the SNR is described in Annex A.

7.10.3.10.2 Structure

The semantics of this primitive are as follows:

PHY_SNR.get {}

7.10.3.10.3 Use

The primitive is generated by the MAC layer when it needs to know the SNR in order to analyse channel characteristics and invoke management procedures for robustness, if required.

7.10.3.11 PHY_SNR.confirm

7.10.3.11.1 Function

The PHY_SNR.confirm primitive is passed to the MAC layer entity by the PHY layer entity in response to a PHY_SNR.get command.

7.10.3.11.2 Structure

The semantics of this primitive are as follows:

PHY_SNR.confirm {*SNR*}

The *SNR* parameter refers to the signal-to-noise ratio, defined as the ratio of the measured received signal level to noise level of the last received PPDU. It may take one of eight values. The mapping of the 3-bit index to the actual SNR value, as calculated in Annex A is given below:

0: ≤ 0 dB

1: ≤ 3 dB

2: ≤ 6 dB

...

7: > 18 dB

7.10.3.12 PHY_ZCT.get

7.10.3.12.1 Function

The PHY_ZCT.get primitive is passed to the PHY layer entity by the MAC layer entity to get the zero cross time of the mains and the time between the last transmission or reception and the zero cross of the mains.

7.10.3.12.2 Structure

The semantics of this primitive are as follows:

PHY_ZCT.get {}

7.10.3.12.3 Use

The primitive is generated by the MAC layer when it needs to know the zero cross time of the mains, e.g., in order to calculate the phase to which the node is connected.

7.10.3.13 PHY_ZCT.confirm

7.10.3.13.1 Function

The PHY_ZCT.confirm primitive is passed to the MAC layer entity by the PHY layer entity in response to a PHY_ZCT.get command.

7.10.3.13.2 Structure

The semantics of this primitive are as follows:

PHY_ZCT.confirm {*Time*}.

The *Time* parameter is the instant in time at which the last zero-cross event took place.

7.10.4 PHY management primitives

PHY layer management primitives enable the PHY layer to interface to the MAC layer. Implementation of these primitives is optional. Refer to Figure 7-14 for the general structure of the PHY layer management primitives.

Table 7-4 – PHY layer management primitives

Primitive	set	get	confirm
PLME_RESET	X		X
PLME_SLEEP	X		X
PLME_RESUME	X		X
PLME_TESTMODE	X		X
PLME_GET		X	X

7.10.4.1 PLME_RESET.request

7.10.4.1.1 Function

The PLME_RESET.request primitive is invoked to request the PHY layer to reset its present functional state. As a result of this primitive, the PHY should reset all internal states and flush all buffers to clear any queued receive or transmit data. All the SET primitives are invoked by the PLME, and addressed to the PHY to set parameters in the PHY. The GET primitive is also sourced by the PLME, but is used only to read PHY parameters.

7.10.4.1.2 Structure

The semantics of this primitive are as follows:

PLME_RESET.request{}

7.10.4.1.3 Use

The upper layer management entities will invoke this primitive to tackle any system level anomalies that require aborting any queued transmissions and restart all operations from the initialization state.

7.10.4.2 PLME_RESET.confirm

7.10.4.2.1 Function

The PLME_RESET.confirm is generated in response to a corresponding PLME_RESET.request primitive. It provides indication if the requested reset was performed successfully or not.

7.10.4.2.2 Structure

The semantics of this primitive are as follows:

PLME_RESET.confirm{*Result*}.

The *Result* parameter shall have one of the following values:

0: Success

1: Failure (the requested reset failed due to internal implementation issues).

7.10.4.2.3 Use

The primitive is generated in response to a PLME_RESET.request.

7.10.4.3 PLME_SLEEP.request

7.10.4.3.1 Function

The PLME_SLEEP.request primitive is invoked to request the PHY layer to suspend its present activities including all reception functions. The PHY layer should complete any pending transmission before entering into a sleep state.

7.10.4.3.2 Structure

The semantics of this primitive are as follows:

PLME_SLEEP.request{}

7.10.4.3.3 Use

This primitive is designed to help optimize power consumption.

7.10.4.4 PLME_SLEEP.confirm

7.10.4.4.1 Function

The PLME_SLEEP.confirm is generated in response to a corresponding PLME_SLEEP.request primitive and provides information if the requested sleep state has been entered successfully or not.

7.10.4.4.2 Structure

The semantics of this primitive are as follows:

PLME_SLEEP.confirm{*Result*}.

The *Result* parameter shall have one of the following values:

- 0: success
- 1: failure (the requested sleep failed due to internal implementation issues)
- 2: PHY layer is already in sleep state.

7.10.4.4.3 Use

The primitive is generated in response to a PLME_SLEEP.request

7.10.4.5 PLME_RESUME.request

7.10.4.5.1 Function

The PLME_RESUME.request primitive is invoked to request the PHY layer to resume its suspended activities. As a result of this primitive, the PHY layer shall start its normal transmission and reception functions.

7.10.4.5.2 Structure

The semantics of this primitive are as follows:

PLME_RESUME.request{}

7.10.4.5.3 Use

This primitive is invoked by upper layer management entities to resume normal PHY layer operations, assuming that the PHY layer is presently in a suspended state as a result of previous PLME_SLEEP.request primitive.

7.10.4.6 PLME_RESUME.confirm

7.10.4.6.1 Function

The PLME_RESUME.confirm is generated in response to a corresponding PLME_RESUME.request primitive and provides information about the requested resumption status.

7.10.4.6.2 Structure

The semantics of this primitive are as follows:

PLME_RESUME.confirm{*Result*}

The *Result* parameter shall have one of the following values:

- 0: success
- 1: failure (the requested resume failed due to internal implementation issues)
- 2: PHY layer is already in fully functional state.

7.10.4.6.3 Use

The primitive is generated in response to a PLME_RESUME.request

7.10.4.7 PLME_TESTMODE.request

7.10.4.7.1 Function

The PLME_TESTMODE.request primitive is invoked to enter the PHY layer into test mode (specified by the mode parameter). A specific functional mode out of the various possible modes is provided as an input parameter. Following the receipt of this primitive, the PHY layer should complete any pending transmissions in its buffer before entering the requested test mode.

7.10.4.7.2 Structure

The semantics of this primitive are as follows:

PLME_TESTMODE.request{*enable, mode, modulation, pwr_level*}.

The *enable* parameter starts or stops the test mode and may take one of two values:

- 0: stop test mode and return to normal functional state
- 1: transit from present functional state to test mode.

The *mode* parameter enumerates specific functional behaviours to be exhibited while the PHY is in test mode. It may have either of the two values:

- 0: continuous transmit
- 1: transmit with 50% duty cycle.

The *modulation* parameter specifies which modulation scheme is used during transmissions. It may take any of the following 8 values:

- 0: DBPSK
- 1: DQPSK
- 2: D8PSK
- 3: not used
- 4: DBPSK + convolutional code
- 5: DQPSK + convolutional code
- 6: D8PSK + convolutional code
- 7: not used

The *pwr_level* parameter specifies the relative level at which the test signal is transmitted. It may take either of the following values:

- 0: maximal output level (MOL)
- 1: MOL – 3 dB
- 2: MOL – 6 dB
- ...
- 7: MOL – 21 dB

7.10.4.7.3 Use

This primitive is invoked by a management entity when specific tests are required to be performed.

7.10.4.8 PLME_TESTMODE.confirm

7.10.4.8.1 Function

The PLME_TESTMODE.confirm is generated in response to a corresponding PLME_TESTMODE.request primitive to indicate if the transition to "Testmode" was successful or not.

7.10.4.8.2 Structure

The semantics of this primitive are as follows:

PLME_TESTMODE.confirm{*Result*}.

The *Result* parameter shall have one of the following values:

- 0: success
- 1: failure (transition to "Testmode" failed due to internal implementation issues)
- 2: PHY layer is already in "Testmode"

7.10.4.8.3 Use

The primitive is generated in response to a PLME_TESTMODE.request.

7.10.4.9 PLME_GET.request

7.10.4.9.1 Function

The PLME_GET.request queries information about a given PIB attribute.

7.10.4.9.2 Structure

The semantics of this primitive are as follows:

PLME_GET.request{*PIBAttribute*}

The *PIBAttribute* parameter identifies specific attributes as enumerated in ID fields of tables that enumerate PIB attributes.

7.10.4.9.3 Use

This primitive is invoked by the management entity to query one of the available PIB attributes.

7.10.4.10 PLME_GET.confirm

7.10.4.10.1 Function

The PLME_GET.confirm primitive is generated in response to the corresponding PLME_GET.request primitive.

7.10.4.10.2 Structure

The semantics of this primitive are as follows:

PLME_GET.confirm{*status*, *PIBAttribute*, *PIBAttributeValue*}

The *status* parameter reports the result of requested information and can have one of the values in Table 7-5.

Table 7-5 – Values of the status parameter in PLME_GET.confirm primitive

Result	Description
Done = 0	Parameter read successfully
Failed = 1	Parameter read failed due to internal implementation reasons.
BadAttr = 2	Specified PIBAttribute is not supported

The *PIBAttribute* parameter identifies specific attributes as enumerated in ID fields of tables that enumerate PIB attributes.

The *PIBAttributeValue* parameter specifies the value associated with a given "*PIBAttribute*".

7.10.4.10.3 Use

This primitive is generated by the PHY layer in response to a PLME_GET.request primitive.

8 Data link layer specifications

8.1 Overview

A subnetwork can be logically seen as a tree structure with two types of nodes: the base node and service nodes.

- Base node: it is at the root of the tree structure and acts as a master node that provides all subnetwork elements with connectivity. It manages the subnetwork resources and connections. There is only one base node in a subnetwork. The base node is initially the subnetwork itself and any other node should follow a registration process to enrol itself on the subnetwork.
- Service nodes: they are either leaves or branch points of the tree structure. They are initially in a disconnected functional state and follow the registration process in clause 8.6.1 to become part of the subnetwork. Service nodes have two functions in the subnetwork: keeping connectivity to the subnetwork for their application layers and switching other nodes' data to propagate connectivity.

Devices elements that exhibit base node functionality continue to do so as long as they are not explicitly reconfigured by mechanisms that are beyond the scope of this Recommendation. Service nodes, on the other hand, change their behaviour dynamically from "terminal" functions to "switch" functions and vice-versa. The changing of functional states occurs in response to certain pre-defined events on the network. Figure 8-1 shows the functional state transition diagram of a service node.

The three functional states of a service node are disconnected, terminal and switch.

- Disconnected: is the initial functional state for all service nodes. When disconnected, a service node is not able to communicate data or switch other nodes' data; its main function is to search for a subnetwork within its reach and try to register on it.
- Terminal: when in this functional state a service node is able to establish connections and communicate data, but it is not able to switch other nodes' data.
- Switch: when in this functional state a service node is able to perform all terminal functions. Additionally, it is able to forward data to and from other nodes in the same subnetwork. It is a branch point on the tree structure.

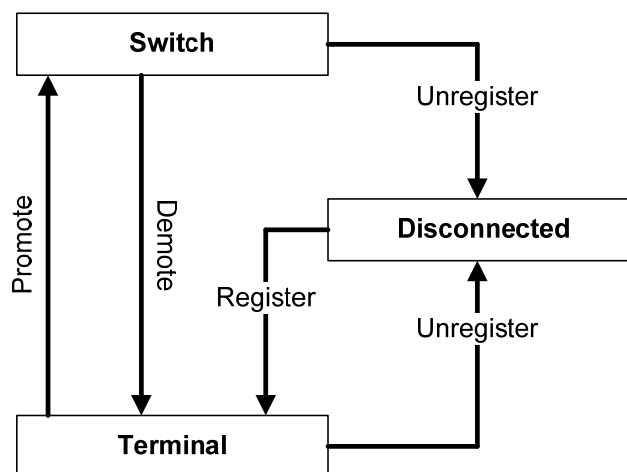


Figure 8-1 – Service node states

The events and associated processes that trigger changes from one functional state to another are given below.

- Registration: the process by which a service node includes itself in the base node's list of registered nodes. Its successful completion means that the service node is part of a subnetwork. Thus, it represents the transition between disconnected and terminal.
- Unregistration: the process by which a service node removes itself from the base node's list of registered nodes. Unregistration may be initiated by either the service node or base node. A service node may unregister itself to find a better point of attachment, i.e., change the switch node through which it is attached to the network. A base node may unregister a registered service node as a result of failure of any of the MAC procedures. Its successful completion means that the service node is disconnected and no longer part of a subnetwork.
- Promotion: the process by which a service node is qualified to switch (repeat, forward) data traffic from other nodes and act as a branch point on the subnetwork tree structure. A successful promotion represents the transition between terminal and switch. When a service node is disconnected it cannot directly transition to switch.
- Demotion: the process by which a service node ceases to be a branch point on the subnetwork tree structure. A successful demotion represents the transition between switch and terminal.

8.2 Addressing

8.2.1 General

Each node has a 48-bit universal MAC address, defined in [IEEE 802-2001] and called EUI-48. Every EUI-48 is assigned during the manufacturing process and it is used to uniquely identify a node during the registration process.

The EUI-48 of the base node uniquely identifies its subnetwork. This EUI-48 is called the subnetwork address (SNA).

The local switch identifier (LSID) is a unique 8-bit identifier for each switch node inside a subnetwork. The subnetwork base node assigns an LSID during the promotion process. A switch node is universally identified by the SNA and LSID. LSID = 0x00 is reserved for the base node. LSID = 0xFF is reserved to mean "unassigned" or "invalid" in certain specific fields (see Table 8-13). This special use of the 0xFF value is always made explicit when describing those fields and it shall not be used in any other field.

During its registration process, every service node receives a 14-bit local node identifier (LNID). The LNID identifies a single service node among all service nodes that directly depend on a given switch. The combination of a service node's LNID and SID (its immediate switch's LSID) forms a 22-bit node identifier (NID). The NID identifies a single service node in a given subnetwork. LNID = 0x0000 cannot be assigned to a terminal, as it refers to its immediate switch. LNID = 0x3FFF is reserved for broadcast and multicast traffic (see clause 8.2.3 for more information). In certain specific fields, the LNID = 0x3FFF may also be used as "unassigned" or "invalid" (see Tables 8-1 and 8-9). This special use of the 0x3FFF value is always made explicit when describing the said fields and it shall not be used in this way in any other field.

During connection establishment a 9-bit local connection identifier (LCID) is reserved. The LCID identifies a single connection in a node. The combination of NID and LCID forms a 31-bit connection identifier (CID). The CID identifies a single connection in a given subnetwork. Any connection is universally identified by the SNA and CID. LCID values are allocated with the following rules:

LCID=0x000 to 0x0FF, for connections requested by the base node. The allocation shall be made by the base node.

LCID=0x100 to 0x1FF, for connections requested by a service node. The allocation shall be made by a service node.

The full addressing structure and field lengths are shown in Figure 8-2.

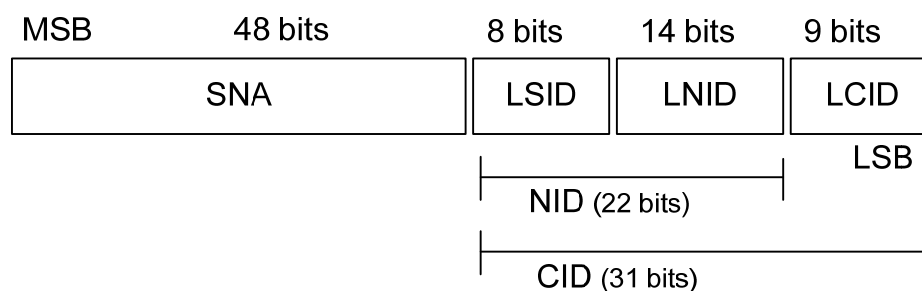


Figure 8-2 – Addressing structure

When a service node in terminal state starts promotion process, the base node allocates a unique switch identifier which is used by this device after transition to the switch state as SID of this switch. The promoted service node continues to use the same NID that it used before promotion, i.e., it maintains the SID of its next level switch for addressing all traffic generated/destined to its local application processes. To maintain distinction between the two switch identifiers, the switch identifier allocated to a service node during its promotion is referred to as a local switch identifier (LSID). Note that the LSID of a switch device will be an SID of devices that connects to the subnetwork through it.

Each service node has a level in the topology tree structure. Service nodes which are directly connected to the base node have level 0. The level of any service node not directly connected to the base node is the level of its immediate switch plus one.

8.2.2 Example of address resolution

Figure 8-3 shows an example where disconnected service nodes are trying to register on the base node. In this example, addressing will have the following nomenclature: (SID, LNID). Initially, the only node with an address is base node A, which has an NID=(0, 0).

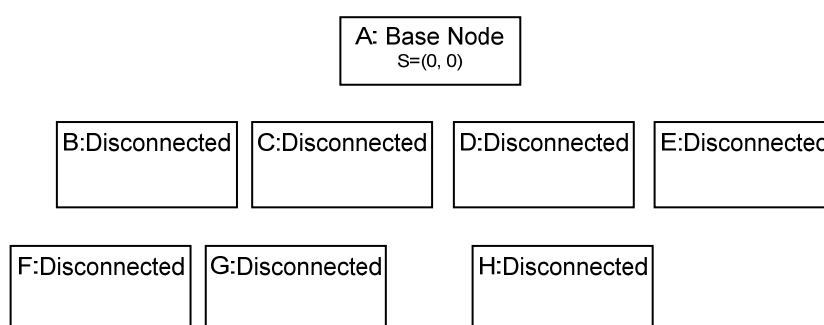


Figure 8-3 – Example of address resolution: phase 1

Every other node of the subnetwork will try to register on the base node. Only B, C, D and E nodes are able to register on this subnetwork and get their NIDs. Figure 8-4 shows the status of nodes after the registration process. Since they have registered on the base node, they get the SID of the base node and a unique LNID. The level of newly registered nodes is 0 because they are connected directly to the base node.

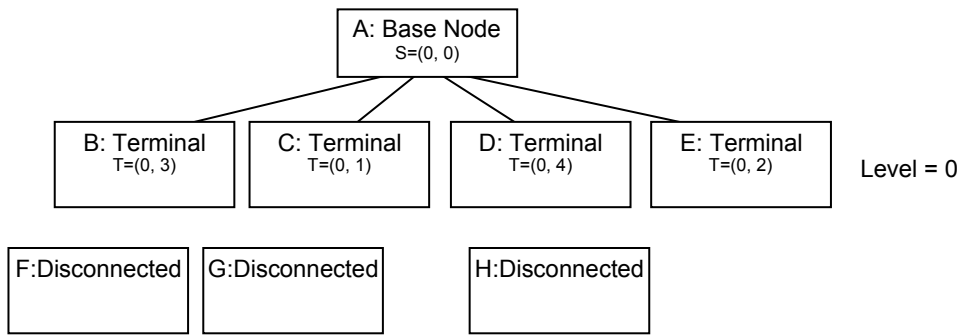


Figure 8-4 – Example of address resolution: phase 2

Nodes F, G and H cannot connect directly to the base node which is currently the only switch in the subnetwork. F, G and H will send promotion needed PDU (PNPDU) broadcast requests which will result in nodes B and D requesting promotion for themselves in order to extend the subnetwork range. During promotion, they will both be assigned unique SIDs. Figure 8-5 shows the new status of the network after the promotion of nodes B and D. Each switch node will still use the NID that was assigned to it during the registration process for its own communication as a terminal node. The new SID shall be used for all switching functions.

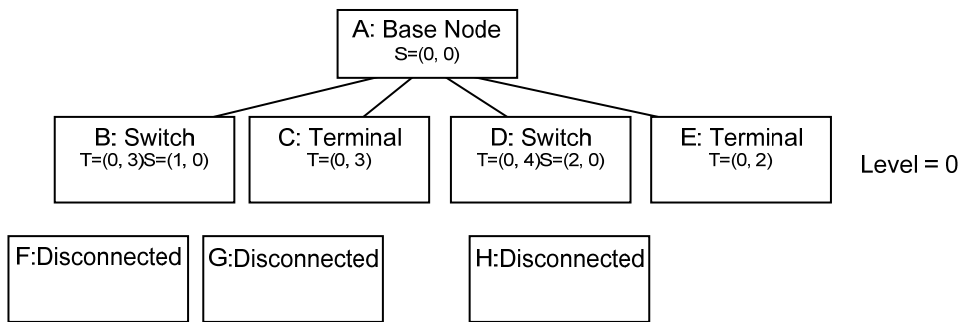


Figure 8-5 – Example of address resolution: phase 3

On completion of the B and D promotion process, nodes F, G and H shall start their registration process and have a unique LNID assigned. Every node on the subnetwork will then have a unique NID to communicate like a terminal, and switch nodes will have unique SIDs for switching purposes. The level of newly registered nodes is 1 because they register with level 0 nodes. On the completion of topology resolution and address allocation, the example subnetwork would be as shown in Figure 8-6.

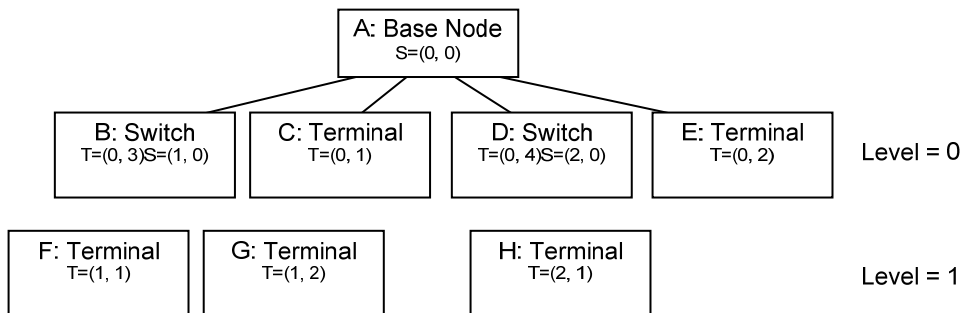


Figure 8-6 – Example of address resolution: phase 4

8.2.3 Broadcast and multicast addressing

Multicast and broadcast addresses are used for communicating data to multiple nodes. There are several broadcast and multicast address types depending on the context associated with the traffic flow. Table 8-1 describes different broadcast and multicast addressing types and the SID and LNID fields associated with each one.

Table 8-1 – Broadcast and multicast address

Type	LNID	Description
Broadcast	0x3FFF	Using this address as a destination, the packets should reach every node of the subnetwork.
Multicast	0x3FFE	This type of address refers to multicast groups. The multicast group is defined by the LCID.
Unicast	not 0x3FFF not 0x3FFE	The address of this type refers to the only node of the subnetwork whose SID and LNID match the address fields.

8.3 MAC functional description

8.3.1 Service node start-up

A service node is initially disconnected. The only functions that may be performed in a disconnected functional state are: reception of any beacons on the channel and sending of the PNPDU. Each service node shall maintain a switch table that is updated with the reception of a beacon from any new switch node. Based on local implementation policies, a service node may select any switch node from the switch table and proceed with the registration process with that switch node. The criteria for selecting a switch node from the switch table is beyond the scope of this Recommendation.

A service node shall listen on the channel for at least *macMinSwitchSearchTime* before deciding that no beacon is being received. It may optionally add some random variation to *macMinSwitchSearchTime*, but this variation cannot be more than 10per cent of *macMinSwitchSearchTime*. If no beacons are received in this time, the service node shall broadcast a PNPDU. The PNPDU shall be broadcast with the most robust modulation scheme to ensure maximum coverage. A service node seeking promotion of any of the terminal nodes in its proximity shall not transmit more than *macMaxPromotionPdu* PNPDU per *macPromotionPduTxPeriod* units of time. The service nodes shall also ensure that the broadcast of PNPDU is randomly spaced. There must always be a random time separation between successive broadcasts.

So as not to flood the network with PNPDU, especially in cases where several devices are powered up at the same time, the terminal nodes shall reduce the PNPDU transmission rate by a factor of PNPDU received from other sources. For example, if a node receives one PNPDU when it is transmitting its own PNPDU, it shall reduce its own transmissions to no more than *macMaxPromotionPdu/2* per *macPromotionPduTxPeriod* units of time. Likewise, if it receives PNPDU from two different sources, it shall slow down its rate to no more than *macMaxPromotionPdu/3* per *macPromotionPduTxPeriod* units of time.

On the selection of a specific switch node, a service node shall start a registration process by transmitting the REG control packet (clause 8.4.5.3) to the base node. The switch node through which the service node intends to carry out its communication is indicated in the REG control packet.

8.3.2 Starting and maintaining subnetworks

Base nodes are primarily responsible for setting up and maintaining a subnetwork. In order to execute the latter, the base node shall perform the following:

- **Beacon transmission:** The base node and all the switch nodes on the subnetwork shall broadcast beacons at fixed intervals of time. The base node shall always transmit exactly one beacon per frame. Switch nodes shall transmit beacons with a frequency prescribed by the base node at the time of their promotion.
- **Promotion and demotion of terminals and switches:** All promotion requests generated by terminal nodes upon receipt of the PNPDU are directed to the base node. The base node maintains a table of all the switch nodes on the subnetwork and allocates a unique SID to new incoming requests. Upon receipt of multiple promotion requests, the base node can, at its own discretion, reject some of the requests. Likewise, the base node is responsible for demoting registered switch nodes. The demotion may either be initiated by the base node (based on an implementation-dependent decision process) or be requested by the switch node itself.
- **Registration management:** The base node receives registration requests from all new nodes trying to be part of the subnetwork it manages. The base node shall process each registration request it receives and respond with an accept or reject message. When the base node accepts the registration of a service node, it shall allocate a unique NID to it to be used for all subsequent communication on the subnetwork. Likewise, the base node is responsible for deregistering any registered service nodes. The unregistration may be initiated by the base node (based on an implementation-dependent decision process) or requested by the service node itself.
- **Connection set-up and management:** The MAC layer specified in this Recommendation is connection-oriented, implying that data exchange is necessarily preceded by connection establishment. The base node is always required for all connections on the subnetwork, either as an end point of the connection or as a facilitator (direct connections; clause 8.3.6) of the connection.
- **Channel access arbitration:** The usage of the channel by devices conforming to this Recommendation may be controlled and contention-free at certain times and open and contention-based at others. The base node prescribes which usage mechanism shall be in force at what time and for how long. Furthermore, the base node shall be responsible for assigning the channel to specific devices during contention-free access periods.
- **Distribution of random sequence for deriving encryption keys.** When using security profile 1 (see clause 8.3.8.2.2), all control messages in this MAC specification shall be encrypted before transmission. Besides control messages, data transfers may be optionally encrypted as well. The encryption key is derived from a 128-bit random sequence. The base node shall periodically generate a new random sequence and distribute it to the entire subnetwork, thus helping to maintain the subnetwork security infrastructure.
- **Multicast group management:** The base node shall maintain all multicast groups on the subnetwork. This shall require the processing of all join and leave requests from any of the service nodes and the creation of unsolicited join and leave messages from base node application requests.

Additional information regarding promotion and connection procedures can be found in clauses 8.6.3 and 8.6.6.

8.3.3 Channel access

8.3.3.1 General

Devices on a subnetwork access the channel based on specific guidelines laid down in this clause. Time is divided into composite units of abstraction for channel usage, called MAC frames. The service nodes and base node on a subnetwork can access the channel in the shared-contention period (SCP) or request a dedicated contention-free period (CFP).

CFP channel access needs devices to request allocation from the base node. Depending on channel usage status, the base node may grant access to the requesting device for a specific duration or deny the request.

SCP channel access does not require any arbitration. However, the transmitting devices need to respect the SCP timing boundaries in a MAC frame. The composition of a MAC frame in terms of SCP and CFP is communicated in every frame as part of the beacon.

A MAC frame is comprised of one or more beacons, one shared-contention period and zero or one contention-free period (CFP). When present, the length of the CFP is indicated in the BPDU.



Figure 8-7 – Structure of a MAC frame

8.3.3.2 Beacon

8.3.3.2.1 General

A BPDU is transmitted by the base node every $(MACFrameLength - MACBeaconLength)$ symbols. The switch nodes also transmit the BPDU to maintain their part of the subnetwork. They transmit BPDUs at regular times, but the transmission frequency does not need to be the same as that of the base node, i.e., a switch node may not transmit its BPDU in every frame.

A beacon is always $MACBeaconLength$ symbols long. This length is the beacon duration excluding the PHY PREAMBLE overhead. Since the BPDU is to be received by all devices in the originating switch domain, it is transmitted with the most robust PHY modulation scheme and FEC coding at the maximum output power level implemented in the device. Details of the BPDU structure and contents are given in clause 8.4.4.

All service nodes shall track beacons as explained in clause 8.3.4.1.

8.3.3.2.2 Beacon slots

A single frame may contain $macBeaconsPerFrame$ BPDUs. The unit of time in which a BPDU is transmitted, is referred to as a beacon slot. All beacon slots are located at the beginning of a frame, as shown in Figure 8-7 above. The first beacon slot in every frame is reserved for the base node. The number of beacon slots in a frame may change from one frame to another and is indicated by the base node in its BPDU.

The switch nodes are allocated a beacon slot at the time of their promotion. Following the PRO control packet, the base node transmits the BSI control packet that would list specific details on which beacon slot should be used by the new switch device.

The number of beacon slots in a frame should be increased from 1 to at least 2 on the promotion of the first switch device on a subnetwork by the base node. Similarly, a base node cannot decrease the number of beacon slots in the subnetwork to 1 when there is a switch node on its subnetwork.

With the registration of each new switch on the subnetwork, the base node may change the beacon slot or BPDU transmission frequency (or both) of already registered switch devices. When such a change occurs, the base node transmits a beacon slot information (BSI) control packet to each individual switch device that is affected. The switch device addressed in the BSI packet sends an acknowledgement back to the base node. Switch devices are required to relay the BSI control packet that is addressed to switch devices connected through them. During the reorganization of beacon slots, if there is a change in the beacon-slot count per frame, the base node should transmit an FRA (FRAME) control packet to the entire subnetwork. The FRA control packet is an indication of change in the overall frame structure. In this specific case, it would imply an increase in SCP slots and a decrease in the number of beacon slots.

Switch devices that receive an FRA control packet should relay it to their entire control domain because FRA packets are broadcast information about changes to frame structures.

This is required for the entire subnetwork to have a common understanding of frame structure, especially in regions where the controlling switch devices transmit BPDUs at frequencies below once per frame.

Figure 8-8 below shows a sample beacon-slot change sequence for an existing switch device. The example shows a beacon-slot change triggered by the promotion of a terminal device (PRO_ACK). In this case, the promotion is followed by a change in both the number of beacon slots per frame and of the specific beacon-slot parameters already allocated to a switch.

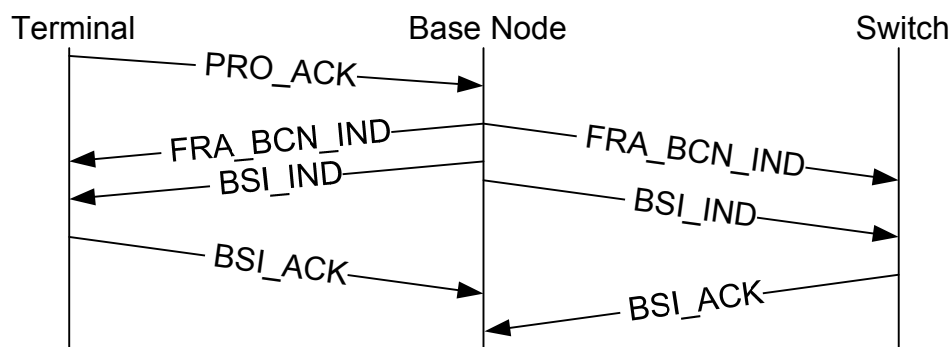


Figure 8-8 – Example of control packet sequencing following a promotion

8.3.3.2.3 Beacon-slot allocation policy

The beacon-slot allocation policy shall ensure that during promotion a service node never receives a BSI control packet that enforces it to transmit a beacon consecutive to every beacon of the node it is registered to.

This behaviour shall be ensured if the BSI information follows one or more of the following rules (BCN represents the information of the beacons the service node is registered to):

- BSI.SLT is not consecutive to BCN.POS
- BSI.SEQ is not equal to any BCN.SEQ in a superframe
- BSI.FRQ is greater than BCN.FRQ.

8.3.3.2.4 Beacon superframes

When changing the frame structure, to add or remove beacon slots or to change which beacon slot a switch should use, it is necessary to indicate when such a change should occur. All nodes must change at the same time otherwise there will be collisions with the beacons etc.

To solve this problem a beacon superframe is defined. Each beacon contains a 5 bit sequence number. Thus 32 frames form a superframe. Any messages which contain changes to the structure or usage of the frame include a sequence number for when the change should occur. The changes requested should only happen when the beacon sequence number matches the sequence number in the change request.

8.3.3.3 Shared-contention period

8.3.3.3.1 General

Shared-contention period (SCP) is the time when any of the devices on the subnetwork can transmit data. The SCP starts immediately after the end of the beacon slot(s) in a frame. Collisions resulting from simultaneous attempt to access the channel are avoided by the CSMA-CA mechanism specified in this clause.

The length of the SCP may change from one frame to another and is indicated by information in the beacon. At all times, the SCP is at least $MACMinSCPLength$ symbols long. The maximum permissible length of an SCP in a frame is $(MACFrameLength - MACBeaconLength)$ symbols. Maximum length SCPs can only occur when there are no dedicated channel access grants to any of the devices (no CFP) on a subnetwork that has no switch nodes (only one beacon slot).

The use of SCP is not restricted to frames in which beacons are received. In lower levels of the subnetwork, the controlling switch node may transmit beacons at a much lower frequency than once per frame. For these parts of the subnetwork, the frame structure would still continue to be the same in frames where no beacons are transmitted. Thus, the service nodes in that segment may still use SCP at their discretion.

8.3.3.3.2 CSMA-CA algorithm

The CSMA-CA algorithm implemented in devices works as shown in Figure 8-9.

Implementations start with a random backoff time ($macSCPRBO$) based on the priority of data queued for transmission. $MACPriorityLevels$ levels of priority need to be defined in each implementation, with a lower value indicating higher priority. In the case of data aggregation, the priority of aggregate bulk is governed by the highest priority data it contains. The $MacSCPRBO$ for a transmission attempt is given below:

$$macSCPRBO = \text{random}(0, \text{MIN}((2^{(\text{Priority} + \text{txAttempts})} + 1), (\text{macSCPLength}/2)))$$

Before a backoff period starts, a device should ensure that the remaining SCP time is long enough to accommodate the backoff, the number of iterations for channel-sensing (based on data priority) and the subsequent data transmission. If this is not the case, backoff should be aborted till the SCP starts in the next frame. Aborted backoffs that start in a subsequent frame should not carry $macSCPRBO$ values of earlier attempts. $macSCPRBO$ values should be regenerated on the resumption of the transmission attempt in the SCP time of the next frame.

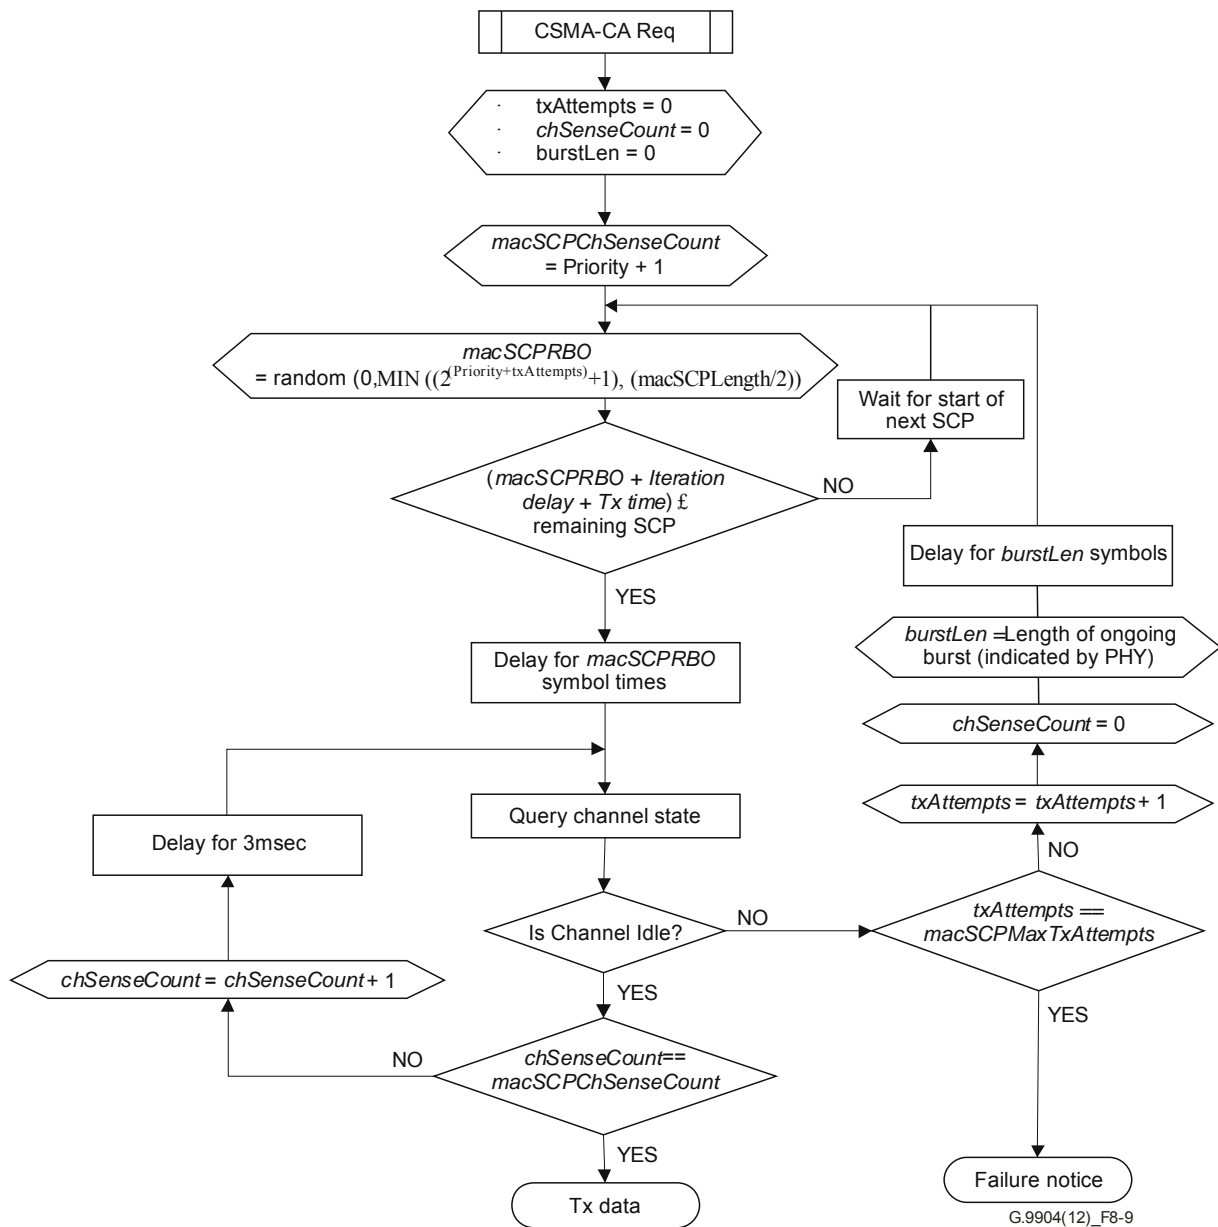


Figure 8-9 – Flow chart for the CSMA-CA algorithm

On completion of the *macSCPRBO* symbol time, implementations perform channel-sensing. Channel sensing shall be performed one or more times depending on the priority of data to be transmitted. The number of times for which an implementation has to perform channel-sensing (*macSCPChSenseCount*) is defined by the priority of the data to be transmitted with the following relation:

$$macSCPChSenseCount = Priority + 1$$

and each channel sense should be separated by a 3 ms delay.

When a channel is sensed to be idle on all *macSCPChSenseCount* occasions, an implementation may conclude that the channel status is idle and carry out its transmission immediately.

During any of the *macSCPChSenseCount* channel-sensing iterations, if the channel is sensed to be occupied, implementations should reset all working variables. The local counter tracking the number of times a channel is found to be busy should be incremented by one and the CSMA-CA process should restart by generating a new *macSCPRBO*. The remaining steps, starting with the backoff, should follow as above.

If the CSMA-CA algorithm restarts *macSCPMaxTxAttempts* number of times due to ongoing transmissions from other devices on the channel, the transmission shall abort by informing the upper layers of CSMA-CA failure.

8.3.3.3.3 MAC control packets

MAC control packets should be transmitted in the SCP with a priority of one. Refer to priorities in clause 8.4.2.3.

8.3.3.4 Contention-free period

Each MAC frame may optionally have a contention-free period where devices are allocated channel time on an explicit request to do so. If no device on a subnetwork requests contention-free channel access, the CFP_ALC_REQ_S may be entirely absent and the MAC frame would only comprise the SCP. All CFP_ALC_REQ_S requests coming from terminal or switch nodes are addressed to the base node. Intermediate switch nodes along the transmission path merely act on the allocation decision by the base node. A single MAC frame may contain up to *MACCFPMaxAlloc* non-overlapping contention-free periods.

Base nodes may allocate overlapping times to multiple requesting service nodes. Such allocations may lead to potential interference. Thus, a base node should make such allocations only when devices that are allocated channel access for concurrent usage are sufficiently separated.

Service nodes make channel allocation requests in a CFP MAC control packet. The base node acts on this request and responds with a request acceptance or denial. In the case of request acceptance, the base node shall respond with the location of allocation time within the MAC frame, the length of allocation time and number of future MAC frames from which the allocation pattern will take effect. The allocation pattern remains effective unless there is an unsolicited location change of the allocation period from the base node (as a result of a channel allocation pattern reorganization) or the requesting service node sends an explicit de-allocation request using a CFP MAC control packet.

Changes resulting from action taken on a CFP MAC control message that impact overall MAC frame structure are broadcast to all devices using an FRA MAC control message.

In a multi-level subnetwork, when a service node that is not directly connected to the base node makes a request for the CFP, the base node shall allocate CFPs to all the intermediate switch nodes so that the entire transit path from the source service node to the base node has contention-free time-slots reserved. The base node shall transmit multiple CFP control packets. The first of these CFP_ALC_IND will be for the requesting service node. Each of the rest will be addressed to an intermediate switch node.

8.3.4 Tracking switches and peers

8.3.4.1 Tracking switches

Service nodes should keep track of all neighbouring switch nodes by maintaining a list of the beacons received. Such tracking shall keep a node updated on reception signal quality from switch nodes other than the one to which it is connected, thus making it possible to change connection points (switch node) to the subnetwork if link quality to the existing point of connectivity degrades beyond an acceptable level.

Note that such a change of point of connectivity may be complex for switch nodes because of devices connected through them. However, at certain times, network dynamics may justify a complex reorganization rather than continue with existing limiting conditions.

8.3.4.2 Tracking disconnected nodes

Terminals shall process all received PNPDU. When a service node is disconnected, it does not have information on current MAC frame structure so the PNPDU may not necessarily arrive during the SCP. Thus, terminals shall also keep track of PNPDU during the CFP or beacon slots.

On processing a received PNPDU, a terminal node may decide to ignore it and not generate any corresponding promotion request (PRO_REQ_S). A terminal node shall ignore no more than *MACMaxPRNIgnore* PNPDU from the same device. Receiving multiple PNPDU from the same device indicates that there is no other device in the vicinity of the disconnected node, implying that there will be no possibility of this new device connecting to any subnetwork if the terminal node does not request promotion for itself.

8.3.5 Switching

8.3.5.1 General

On a subnetwork, the base node cannot communicate with every node directly. Switch nodes relay traffic to/from the base node so that every node on the subnetwork is effectively able to communicate with the base node. Switch nodes selectively forward traffic that originates from or is destined to one of the service nodes in its control hierarchy. All other traffic is discarded by switches, thus reducing traffic flow on the network.

Different names of MAC header and packets are used in this clause. Please refer to clause 8.4.2 to find their complete specification.

8.3.5.2 Switching table

Each switch node maintains a table of other switch nodes that are connected to the subnetwork through it. Maintaining this information is sufficient for switching because traffic to/from the terminal nodes will also contain the identity of their respective switch nodes (PKT.SID). Thus, the switching function is simplified in that maintaining an exhaustive listing of all terminal nodes connected through it is not necessary.

Switch nodes start with no entries in their switching table. The switching table is dynamically updated by keeping track of promotion and demotion control packets flowing on the network. A new entry is created for every promotion acknowledgement (PRO_ACK) that has a PKT.SID matching either the SID of the switch node itself or any of the existing entries in the switching table. Likewise, an entry corresponding to a PRO.NSID field is deleted when a demotion request is acknowledged (PRO_DEM_x).

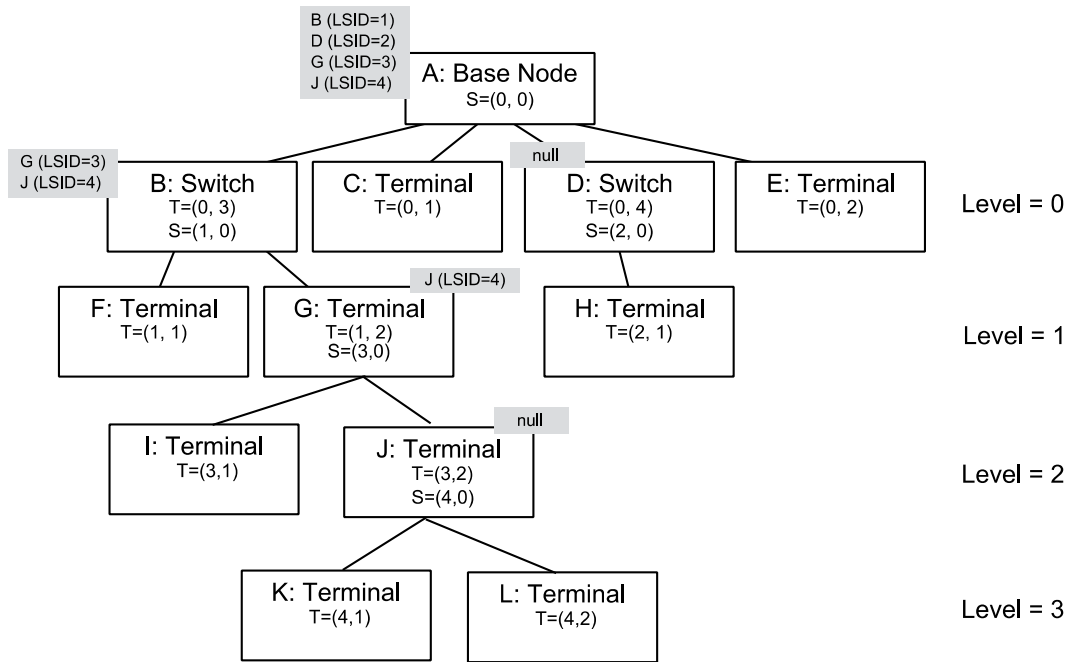


Figure 8-10 – Switching tables example

Figure 8-10 shows an example subnetwork where entries in the switching table of individual switch nodes are highlighted. In this example, when service node G receives a PRO_REQ_B packet for promotion, it turns into a switch node. Its switch identifier will be (PRO.NSID, 0) = (3, 0). The receipt and acceptance of PRO_REQ_B is acknowledged with a PRO_ACK by G. The intermediate switch node B will sniff HDR.DO=0, PKT.CTYPE=3, PKT.SID=1 and PRO.N=0, to conclude that this is a PRO_ACK from one of the service nodes in its own switching hierarchy. Node B will forward this packet towards the base node and it will add PRO.NSID to its switching table, as shown in Figure 8-11.

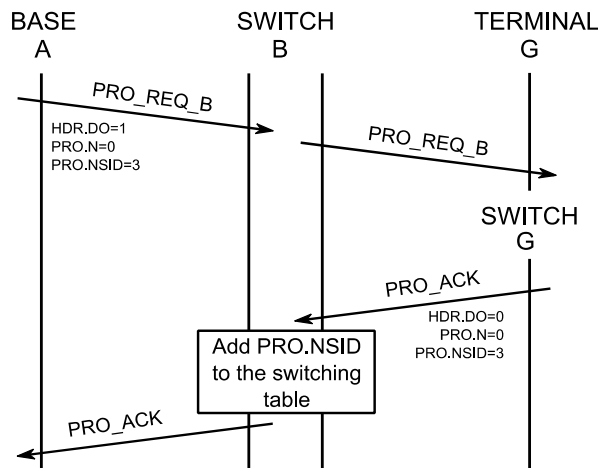


Figure 8-11 – Fill in the switching table

Removing a switch table entry is more complex because of retries. Upon receipt of a demotion acknowledgement (PRO_DEM_x), the switching table entry corresponding to the LSID is marked as to be removed and a timer is started with a value of $((macMaxCtlReTx + 1) \times macCtlReTxTimer)$ seconds. This timer ensures that all retransmit packets which might use the LSID have left the subnetwork. When the timer expires the switch table entry is removed.

8.3.5.3 Switching process

Switch nodes forward traffic to their control domain in a selective manner. The received data shall fulfil the conditions listed below for it to be switched. If the conditions are not met, the data shall be silently discarded.

Downlink packets (HDR.DO=1) shall meet any of the following conditions in order to be switched:

- Destination node of the packet is connected to the subnetwork through this switch node, i.e., PKT.SID is equal to this switch node's SID or its switching table contains an entry for PKT.SID.
- The packet has broadcast destination (PKT.LNID = 0x3FFF) and was sent by the switch this node is registered through (PKT.SID=SID of this switch node).
- The packet has a multicast destination (PKT.LNID=0x3FFE), it was sent by the switch this node is registered through (PKT.SID=SID of this switch node) and at least one of the service nodes connected to the subnetwork through this switch node is a member of the said multicast group, i.e., LCID specifies a group that is requested by any downstream node in its hierarchy.

Uplink packets (HDR.DO=0) shall meet either of the following conditions in order to be switched:

- The packet source node is connected to the subnetwork through this switch node, i.e. PKT.SID is equal to this switch node's SID or its switching table contains an entry for PKT.SID.
- The packet has a broadcast or multicast destination (PKT.LNID = 0x3FFF or 0x3FFE) and was transmitted by a node registered through this switch node (PKT.SID=LSID of this switch node).

If a packet meets previous conditions, it shall be switched. For unicast packets, the only operation to be performed during switching is queueing it to be resent in a MAC PDU with the same HDR.DO.

In the case of broadcast or multicast packets, the PKT.SID must be replaced with:

- the switch node's LSID for downlink packets
- the switch node's SID for uplink packets.

8.3.5.4 Switching of broadcast packets

The switching of broadcast MAC frames operates in a different manner to the switching of unicast MAC frames. Broadcast MAC frames are identified by PKT.LNID=0x3FFF.

When HDR.DO=0, i.e., the packet is an uplink packet, it is unicast to the base node. A switch which receives such a packet should apply the scope rules to ensure that it comes from a lower level and, if so, switch it upwards towards the base. The rules given in clause 8.3.5.3 must be applied.

When HDR.DO=1, i.e., the packet is a downlink packet, it is broadcast to the next level. A switch which receives such a packet should apply the scope rules to ensure that it comes from the higher level and, if so, switch it further to its subnetwork. The most robust PHY modulation scheme and FEC coding at the maximum output power level implemented in the device should be used so that all the devices directly connected to the switch node can receive the packet. The rules given in clause 8.3.5.3 must be applied. The service node should also pass the packet up to its MAC SAP to applications which have registered to receive broadcast packets using the MAC_JOIN service.

When the base node receives a broadcast packet with HDR.DO=0, it should pass the packet up its MAC SAP to applications which have registered to receive broadcast packets. The base node should also transmit the packet as a downlink packet, i.e., HDR.DO=1, using the most robust PHY modulation scheme and FEC coding at the maximum output power level and following the rules given in clause 8.3.5.3.

8.3.5.5 Switching of multicast packets

8.3.5.5.1 General

Multicast packet switching operates in a very similar way to broadcast packet switching. Multicast is an extension of broadcast. If a switching node does not implement multicasting, it should handle all multicast packets as broadcast packets.

Different names of MAC header and packets are used in this clause. Refer to clause 8.4.2 to find proper definitions.

8.3.5.5.2 Multicast switching table

Switch nodes which implement multicast should maintain a multicast switching table. This table contains a list of multicast group LCIDs that have members connected to the subnetwork through the switch node. The LCID of multicast traffic in both downlink and uplink directions is checked for a matching entry in the multicast switching table. Multicast traffic is only switched if an entry corresponding to the LCID is available in the table; otherwise, the traffic is silently discarded.

A multicast switching table is established and managed by examining the multicast join and leave messages (MUL control packet) which pass through the switch. Note that multiple service nodes from a switch node's control hierarchy may be members of the same group.

On a successful group join from a service node in its control hierarchy, a switch node adds a new multicast switch entry for the group LCID, where necessary.

When a successful group leave is indicated, the switch removes the NID from the multicast switch entry. If the multicast switch entry then has no NID associated with it, the multicast switch entry is immediately removed.

switch nodes shall also examine the Keep-Alive packets being passed upwards. When a service node that is also a member of a multicast group fails the Keep-Alive process, its NID is removed from any multicast switch entries and, if necessary, the multicast switch entry is removed.

Switch nodes should use a timer to trigger the actual removal of switch entries. The timer is started when it is decided that an entry should be removed. This timer has the value $((macMaxCtlReTx + 1) \times macCtlReTxTimer)$. Only once the timer has expired is the multicast switch entry removed. This allows the terminal node a short amount of time to flush any remaining multicast packets before the connection is removed and the switch node implementation is simplified since it only needs to process MUL_LEAVE_B or MUL_LEAVE_S (refers to clause 8.4.5.9), but not both.

8.3.5.5.3 Switching process of multicast packets

The multicast packet switching process depends on the packet direction.

When HDR.DO=0 and PKT.LNID=0x3FFE, i.e., the packet is an uplink multicast packet, it is unicast towards the base node. A switch node that receives such a packet should apply the scope rules to ensure it comes from a lower hierarchical level and, if so, switch it upwards towards the base node. No LCID-based filtering is performed. All multicast packets are switched, regardless of any multicast switch entries for the LCID. The coding rate most applicable to the unicast may be used and the rules given in clause 8.3.5.3 shall be applied.

When HDR.DO=1 and PKT.LNID=0x3FFE, i.e., the packet is a downlink multicast packet, the multicast switching table is used. If there is an entry with the LCID corresponding to PKT.LCID in the packet, the packet is switched downwards to the part of subnetwork controlled by this switch. The most robust PHY modulation scheme and FEC coding at the maximum output power level should be used so that all its devices in the lower level can receive the packet. The rules given in clause 8.3.5.3 shall be applied. If the service node is also a member of the multicast group, it should also pass the packet up its MAC SAP to applications which have registered to receive the multicast packets for that group.

When the base node receives a multicast packet with HDR.DO=0 and it is a member of the multicast group, it should pass the packet up its MAC SAP to applications which have registered to receive multicast packets for that group. The base node should switch the multicast packet if there is an appropriate entry in its multicast switching table for the LCID, transmitting the packet as a downlink packet, i.e., HDR.DO=1, using the most robust PHY modulation scheme and FEC coding at the maximum output power level. The rules given in clause 8.3.5.3 shall be used.

8.3.6 Direct connections

8.3.6.1 Direct connection establishment

The direct connection establishment is a little different from a normal connection although the same packets and processes are used. It is different because the initial connection request may not be acknowledged until it is already acknowledged by the target node. It is also different because the CON_REQ_B packets shall carry information for the "direct switch" to update the "direct switching table".

A direct switch is not different from a general switch. It is only a logical distinction of identifying the first common switch between two service nodes that need to communicate with each other. Note that in the absence of such a common switch, the base node would be the direct switch.

There are two different scenarios for using directed connections. These scenarios use the network shown in Figure 8-12.

The first is when the source node does not know the destination service node's EUI-48 address. The service node initiates a connection to the base node and the base node convergence layer redirects the connection to the correct service node.

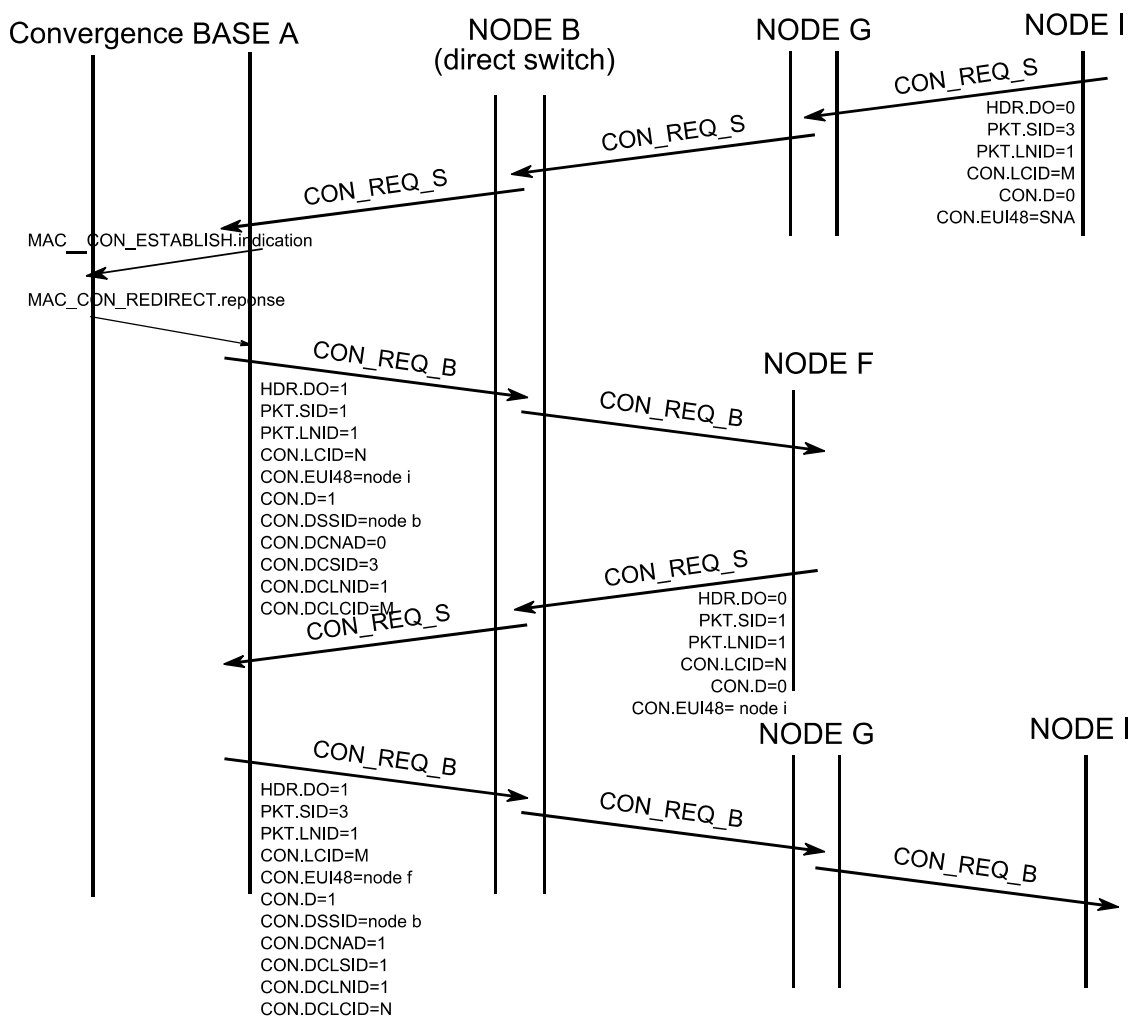


Figure 8-12 – Directed connection to an unknown service node

The steps to establish a direct connection, as shown in Figure 8-12, shall be:

- When node I tries to establish a connection with node F, it shall send a normal connection request (CON_REQ_S).
- Then, due to the fact that the base node knows that F is the target service node, it should send a connection request to F (CON_REQ_B). This packet will carry information for direct switch B to include the connection in its direct switching table.
- F may accept the connection. (CON_REQ_S).
- Now that the connection with F is fully established, the base node will accept the connection with I (CON_REQ_B). This packet will carry information for the direct switch B to include in its direct switching table.

After finishing this connection-establishment process, the direct switch (node B) should contain a direct switching table with the entries shown in Table 8-2.

Table 8-2 – Direct connection example: node B's direct switching table

Uplink			Downlink			
SID	LNID	LCID	DSID	DLNID	DLCID	NAD
1	1	N	3	1	M	0
3	1	M	1	1	N	1

The direct switching table should be updated every time a switch receives a control packet that meets the following requirements:

- it is CON_REQ_B packet: HDR.DO=1, CON.TYPE=1 and CON.N=0
- it contains "direct" information: CON.D=1
- the direct information is for itself: CON.DSSID is the SID of the switch itself.

Then, the direct switching table is updated with the information:

- uplink (SID, LNID, LCID) = (PKT.SID, PKT.LNID, CON.LCID);
- downlink (SID, LNID, LCID, NAD) = (CON.DCSID, CON.DCLNID, CON.DCLCID, CON.DCNAD).

The connection closing packets should be used to remove the entries.

The second scenario for using directed connections is when the initiating service node already knows the destination service node's EUI-48 address. In this case, rather than using the base node's address, it uses the service node's address. In this case, the base node convergence layer is not involved. The base node MAC layer connects service node I directly to service node F. The resulting switch table entries are identical to the previous example. The exchange of signals is shown in Figure 8-13.

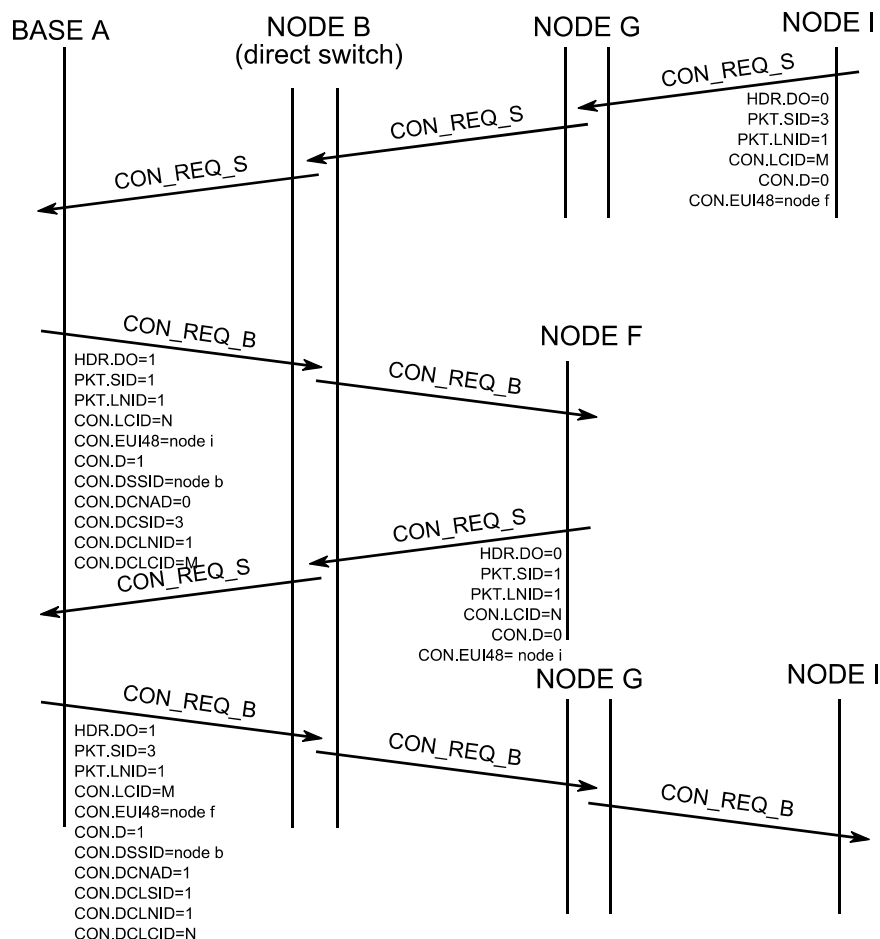


Figure 8-13 – Example of direct connection: connection establishment to a known service node

8.3.6.2 Direct connection release

The release of a direct connection is shown in Figure 8-14. The signalling is very similar to connection establishment for a direct connection. The D fields are used to tell the direct switch which entries it should remove. The direct switching table should be updated every time a switch receives a control packet that meets the following requirements:

- it is CON_CLOSE_B packet: HDR.DO=1, CON.TYPE=1 and CON.N=1
- it contains "direct" information: CON.D=1
- the direct information is for itself: CON.DSSID is the SID of the switch itself.

Then, the direct switching table entry with the following information is removed:

- uplink (SID, LNID, LCID) = (PKT.SID, PKT.LNID, CON.LCID);
- downlink (SID, LNID, LCID, NAD) = (CON.DCSID, CON.DCLNID, CON.DCLCID, CON.DCNAD).

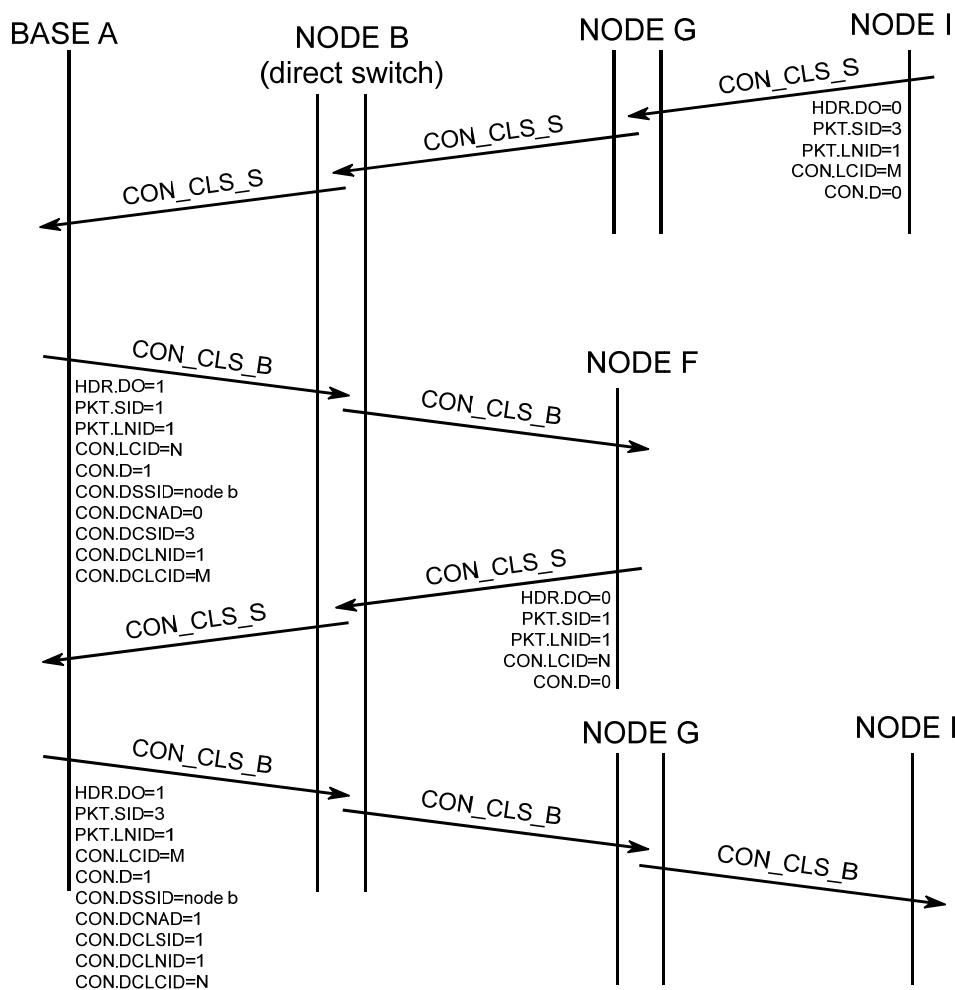


Figure 8-14 – Release of a direct connection

8.3.6.3 Direct connection switching

As explained in clause 8.3.5.3, the normal switching mechanism is intended to be used for forwarding communication data between the base node and each service node. The "direct switching" is a mechanism to let two nodes communicate with each other, switching the packets in a local way, i.e., without passing through the base node. It is not a different form of packet-switching, but rather an additional feature of the general switching process.

The first shared switch in the paths that go from two service nodes to the base node will be called the "direct switch" for the connections between the said nodes. This is the switch that will have the possibility of performing the direct switching to make the two nodes communicate efficiently. As a special case, every switch is the "direct switch" between itself and any node that is lower down in the hierarchy.

The "direct switching table" is a table every switch should contain in order to perform direct switching. Each entry on this table is a direct connection that must be switched directly. It is represented by the origin CID and the destination CID of the direct connection. It is not a record of every connection identifier lower down in its hierarchy, but contains only those that should be directly switched by it. The destination node's ability to receive aggregated packets shall also be included in the "direct switching table" in order to fill the PKT.NAD field.

8.3.6.4 Direct switching operation

If a switch receives an uplink (HDR.DO=0) MAC frame that is to be switched (see clause 8.3.5.3 for the requirements) and its address is in the direct switching table, then the procedure is as follows:

- change the (SID, LNID, LCID, NAD) by the downlink part of the entry in the direct switching table;
- queue the packet to be transmitted as a downlink packet (HDR.DO=1).

8.3.7 Packet aggregation

8.3.7.1 General

The GPDU may contain one or more packets. The functionality of including multiple packets in a GPDU is called packet aggregation. Packet aggregation is an optional part of this Recommendation and devices do not need to implement it for compliance with this Recommendation. It is however suggested that devices should implement packet aggregation in order to improve MAC efficiency.

To maintain compatibility between devices that implement packet aggregation and ones that do not, there must be a guarantee that no aggregation takes place for packets whose data transit path from/to the base node crosses (an) intermediate service node(s) that do(es) not implement this function. Information about the aggregation capability of the data transit path is exchanged during the registration process (8.6.1). A registering service node notifies this capability to the base node in the REG.CAP_PA field (1 bit, see Table 8-8) of its REG_REQ message. It gets feedback from the base node on the aggregation capability of the whole Downlink transit path in the REG.CAP_PA field of the REG_RSP message.

Based on initial information exchanged on registration, each subsequent data packet in either direction contains aggregation information in the PKT.NAD field. In the Downlink direction, the base node will be responsible for filling PKT.NAD based on the value it communicated to the destination service node in the REG.CAP_PA field of the REG_RSP message. Likewise, for uplink data, the source service node will fill PKT.NAD based on the REG.CAP_PA field received in the initial REG_RSP from the base node. The last switch shall use the PKT.NAD field to avoid packet aggregation when forwarding the packet to destination service nodes without packet aggregation capability. Intermediate switch nodes should have information about the aggregation capability in their switching table and shall not aggregate packets when it is known that next level switch node does not support this feature.

Devices that implement packet aggregation shall ensure that the size of the MSDU comprising the aggregates does not exceed the maximum capacity of the most robust transmission scheme of a PHY burst. The most robust transmission scheme refers to the most robust combination of modulation scheme and convolutional coding.

8.3.7.2 Packet aggregation when switching

Switch nodes maintain information on the packet aggregation capability of all entries in their switching table, i.e., of all switches that are connected to the subnetwork through them. This capability information is then used during traffic switching to/from the connected switch nodes.

The packet aggregation capability of a connecting switch node is registered at each transit switch node at the time of its promotion by sniffing relevant information in the PRO_ACK message.

- If the PKT.SID in a PRO_ACK message is the same as the switching node, the node being promoted is connected directly to the said switch node. The aggregation capability of this new switch node is registered as the same as indicated in PKT.NAD of the PRO_ACK packet.
- If the PKT.SID in a PRO_ACK message is different from the SID of the switching node, it implies that the node being promoted is indirectly connected to this switch. The aggregation capability for this new switch node will thus be the same as the aggregation capability registered for its immediate switch, i.e., PKT.SID.

Aggregation while switching packets in the uplink direction is performed if the node performing the switch knows that its uplink path is capable of handling aggregated packets, based on capability information exchanged during registration (REG.CAP_PA field in REG_RSP message).

Downlink packets are aggregated by analysing the following:

- If the PKT.SID is the same as the switching node, then it is the last switching level and the packet will arrive at its destination. In this case, the packet may be aggregated if PKT.NAD=0.
- If the PKT.SID is different, this is not the last level and another switch will receive the packet. The information of whether or not the packet could be aggregated should be extracted from the switching table.

8.3.8 Security

8.3.8.1 General

The security functionality provides the MAC layer with privacy, authentication and data integrity through a secure connection method and a key management policy. All packets must use the negotiated security profile. The only exceptions to this rule are the REG and SEC control messages, and the BPDU and PNPDU PDUs which are transferred non-encrypted.

8.3.8.2 Security profiles

Several security profiles are provided for managing different security needs, which can arise in different network environments. This version of the Recommendation lists two security profiles and leaves scope for adding up to two new security profiles in future versions.

8.3.8.2.1 Security profile 0

Communications which have the security profile 0 are based on the transmission of MAC SDUs without encryption. This profile may be used in application scenarios where either sufficient security is provided by upper communication layers or where security is not a major requirement for application use-case.

8.3.8.2.2 Security profile 1

8.3.8.2.2.1 General

Security profile 1 is based on 128-bit AES encryption of data and its associated CRC. This profile is specified with the aim of fulfilling all security requirements:

- Privacy is guaranteed by the encryption itself and by the fact that the encryption key is kept secret.
- Authentication is guaranteed by the fact that each node has its own secret key known only by the node itself and the base node.
- Data integrity is guaranteed by the fact that the payload CRC is encrypted.

8.3.8.2.2.2 Cryptographic primitives

The cryptographic algorithm used in this Recommendation is the AES, as specified in [PUB 197]. The specification describes the algorithm with three possible key sizes; the 128-bit secret key represents a good level of security for preserving privacy up to 2030 and beyond, as specified in [SP 800-57], page 66, Table 4.

AES is used according to the so-called ECB, as specified in [SP 800-38A]. It is a block-ciphering mode where plain text is divided into 128-bit blocks. Padding is applied if the last block is smaller than 128 bits. Padding is implemented with the addition of a bit equal to 1 and as many zeroes as necessary to reach a length of the string to be encrypted as a multiple of 128 bits. Encryption is performed one block at a time using the same working key for all the data.

8.3.8.2.2.3 Key derivation algorithm

The method for deriving working keys from secret keys is to apply the AES algorithm to a constant (C) as plain text and generation key (GK) as an encryption key. If the constant is shorter than 128 bits, it must be aligned to the LSB, as shown in Figure 8-15. The various key derivation equations specified in the following clauses follow the convention:

Generated Key = AES_enc (*Generation Key*, *Constant*)

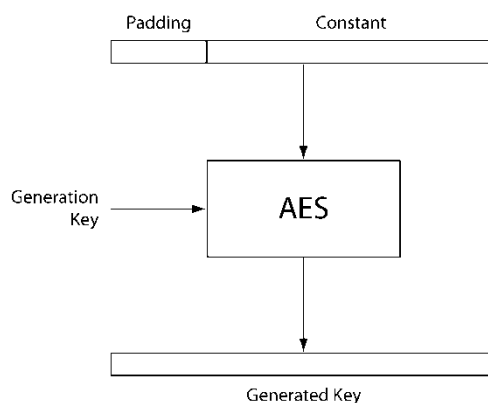


Figure 8-15 – Key derivation concept

8.3.8.3 Negotiation of the security profile

All MAC data, including signalling PDUs (all MAC control packets defined in clause 8.4.5) use the same security profile. This profile is negotiated during the device registration. In the REG_REQ message the terminal indicates a security profile it is able to support in the field REG.SPC. The base node may accept this security profile and so accept the registration, sending back a REG_RSP with the same REG.SPC value. The base node may also accept the registration, however it sets REG.SPC to 0 indicating that security profile 0 is to be used. Alternatively, the base node may reject the registration if the terminal does not provide an acceptable security profile.

It is recommended that the terminal first attempts to register using the highest security profile it supports and only use lower security profiles when the base node rejects the registration request.

8.3.8.4 Key hierarchy

8.3.8.4.1 Security profile 0

Not applicable.

8.3.8.4.2 Security profile 1

Service nodes and base nodes use a set of three working keys to encrypt all data. The keys and their respective usage are:

Initial working key (WK0): This key has limited scope and is used to decrypt the REG.SNK and REG.AUK fields of the REG_RSP message. The WK0 is thus used by a service node in a disconnected functional state. This key is computed using the following formula:

$$WK0 = \text{AES_enc}(USK, 0)$$

Working key (WK): This key is used to encrypt all the unicast data that is transmitted from the base node to a service node and vice versa. Each registered service node would have a unique WK that is known only to the base node and itself. The WK is computed as follows:

$$WK = \text{AES_enc}(USK, \text{Random sequence received in SEC.RAN})$$

Subnetwork working key (SWK): The SWK is shared by the entire subnetwork. To ensure the security of this key, it is never transmitted over the physical channel, but is computed from other keys which are transmitted encrypted in REG and non-encrypted in SEC control packets. The SWK shall be used to encrypt the following:

- broadcast data, including MAC broadcast control packets
- multicast data
- unicast data that is transacted over direct connections, i.e., not involving the base node.

The SWK is computed as follows:

$$SWK = \text{AES_enc}(SNK, \text{Random sequence received in SEC.SNK})$$

The WK and the SWK have a limited validity time related to the random sequence generation period. The random sequence is regenerated and distributed by the base node at least every *MACRandSeqChgTime* seconds through the SEC control packet. If a device does not receive an update of a random sequence within $2 \times \text{MACRandSeqChgTime}$, it should consider the WK and SWK as no longer valid. Lack of availability of WK will render service node to very limited functionality of unregistering from the subnetwork. It is therefore advised that in such cases, service nodes should unregister from the network and initiate a re-registration procedure.

The key derivation procedures have been designed to be indirect and multi-staged to ensure security. The parameters involved in the derivation of the working keys are defined below.

Master key (MK1, MK2). Two master keys (MK1 and MK2) are defined in this Recommendation. MK1 is used to compute the DSK. MK2 is used to compute the KDIV. Both of these keys are administered on the base node by implementation-dependent means that are beyond the scope of this Recommendation. Specifying two master keys makes the USK generation a two stage process, i.e., derivation of DSK and KDIV in the first stage and using them to derive the USK in the second stage. Note that the DSK and KDIV are unique to each registering service node.

Device secret key (DSK). DSK is unique to each service node on the subnetwork and is hard-coded in the device during production. The DSK is constant for the entire life of the service node. The base node uses MK1 to derive service node-specific DSK using the following equation:

$$DSK = \text{AES_enc}(MK1, UI)$$

Key diversifier (KDIV). This quantity is also unique to each service node, but unlike DSK, it does not have to be a fixed constant for the entire life of the service node. The KDIV is provisioned on each service node by means that are beyond the scope of this Recommendation. The base node computes device-specific KDIV using the equation:

$$KDIV = AES_enc (MK2, UI)$$

Unique secret key (USK). The USK is used to derive WK0 and WK as defined in the above equations. The USK is in turn computed by applying AES to KDIV, using DSK as the generation key, as shown in the equation below. Note that this is a single-step process in service nodes because both KDIV and DSK are already known or provisioned, but a three-step process in the base node. The first two steps in the base node comprise deriving the DSK and KDIV using the MK1 and MK2, respectively.

$$USK = AES_enc (DSK, KDIV)$$

Unique identifier (UI): The UI of a service node shall be its EUI-48.

8.3.8.5 Key distribution and management

The security profile for data traffic is negotiated when a device is registered. The REG control packet contains specific fields to indicate the security profile for respective devices. All connections to/from the device would be required to follow the security profile negotiated at the time of registration. There cannot be a difference in security profiles across multiple connections involving the same device. The only exception to this would be the base node.

The SWK used as a working key for non-unicast traffic and direct connections is never transmitted in non-encrypted form over the physical channel. The SEC broadcast messages transmitted by the base node (and relayed by all switch nodes) at regular intervals contain random keys for both unicast and non-unicast traffic. When a device initially registers on a subnetwork, the REG response from the base node contains the random sequence used to derive WK for unicast traffic. The REG message is followed by a unicast SEC message from the base node to the registering device.

8.3.8.6 Encryption

8.3.8.6.1 Security profile 0

Not applicable.

8.3.8.6.2 Security profile 1

Connections working with "Security profile 1" would always transmit a CRC with every packet. This field shall be called SCRC (security CRC) and is calculated over the unencrypted packet payload. The SCRC helps confirm the integrity of the packet on its decryption at the receiving end.

The SCRC shall be calculated as the remainder of the division (Modulo 2) by the generator polynomial $g(x)=x^8+x^2+x+1$ of the polynomial x_8 multiplied by the unencrypted packet payload.

The data block obtained by the concatenation of the unencrypted payload of the packet and the calculated SCRC is padded with a 1 followed by as many zeroes as necessary to reach a multiple of 128 and then divided into 128-bit blocks. The 1 inserted as the first padding bit is useful to detect the start of the padding at the receiver without notification of the number of padded bits.

Each 128-bit block is encrypted with the AES algorithm using a valid working key. The result of this encryption process is the encrypted payload of the packet.

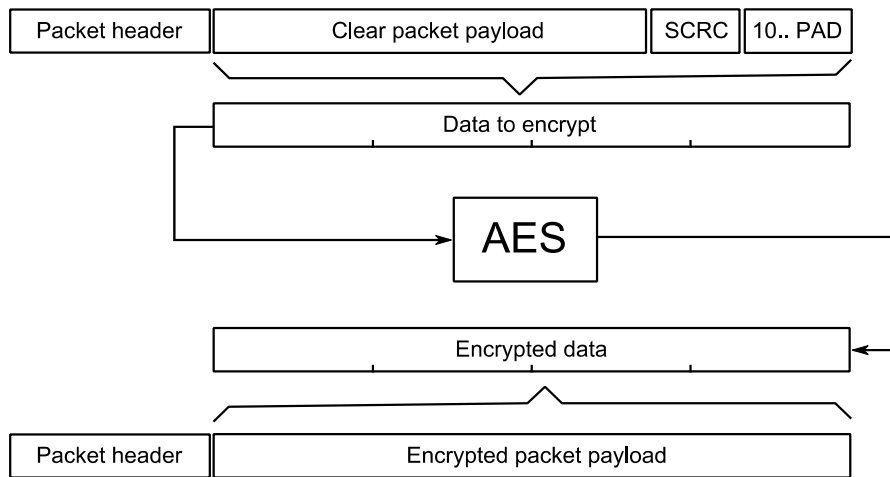


Figure 8-16 – Security profile 1 encryption algorithm

8.4 MAC PDU format

8.4.1 General

There are different types of MAC PDUs for different purposes.

8.4.2 Generic MAC PDU

8.4.2.1 General

Most subnetwork traffic comprises generic MAC PDUs (GPDU). GPDUs are used for all data traffic and most control traffic. All MAC control packets are transmitted as GPDUs.

GPDU composition is shown in Figure 8-17. It is composed of a generic MAC header followed by one or more MAC packets and a 32 bit CRC appended at the end.




Figure 8-17 – Generic MAC PDU format

8.4.2.2 Generic MAC header

The generic MAC header format is represented in Table 8-3. The size of the generic MAC header is 3 bytes. Table 8-3 enumerates each field of a generic MAC header.

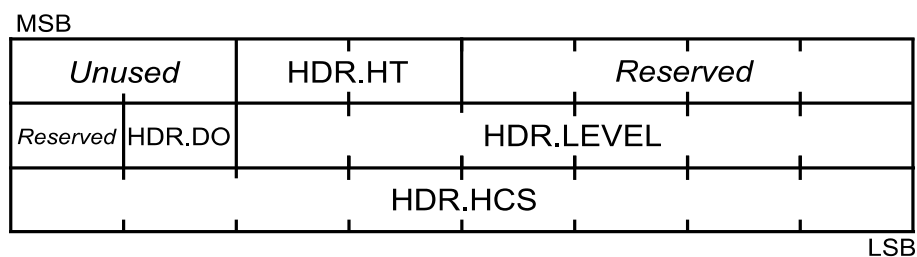


Figure 8-18 – Generic MAC header

Table 8-3 – Generic MAC header fields

Name	Length	Description
<i>Unused</i>	2 bits	Unused bits which are always 0; included for alignment with MAC_H field in the PPDU header (see clause 7.4.3).
HDR.HT	2 bits	Header type HDR.HT = 0 for GPDU
<i>Reserved</i>	5 bits	Always 0 for this version of the Recommendation. Reserved by ITU-T.
HDR.DO	1 bit	Downlink/uplink <ul style="list-style-type: none"> • HDR.DO=1 if the MAC PDU is downlink • HDR.DO=0 if the MAC PDU is uplink
HDR.LEVEL	6 bits	Level of the PDU in switching hierarchy The packets between the level 0 and the base node are of HDR.LEVEL=0. The packets between levels k and k-1 are of HDR.LEVEL=k. <ul style="list-style-type: none"> • If HDR.DO=0, HDR.LEVEL represents the level of the transmitter of this packet. • If HDR.DO=1, HDR.LEVEL represents the level of the receiver of this packet.
HDR.HCS	8 bits	Header check sequence A field for detecting errors in the header and checking that this MAC PDU is from this subnetwork. The transmitter shall calculate the CRC of the SNA concatenated with the first 2 bytes of the header and insert the result into the HDR.HCS field (the last byte of the header). The CRC shall be calculated as the remainder of the division (Modulo 2) of the polynomial $M(x) \cdot x^8$ by the generator polynomial $g(x)=x^8+x^2+x+1$. $M(x)$ is the input polynomial, which is formed by the bit sequence of the concatenation of the SNA and the header excluding the HDR.HCS field, and the MSB of the bit sequence is the coefficient of the highest order of $M(x)$.

8.4.2.3 Packet structure

A packet is comprised of a packet header and packet payload. Figure 8-19 shows the structure.




Figure 8-19 – Packet structure

Packet header is 6 bytes in length and its composition is shown in Figure 8-20. Table 8-4 enumerates the description of each field.




Figure 8-20 – Packet header

To simplify, the text contains references to the PKT.NID fields as the composition of the PKT.SID and PKT.LNID. The field PKT.CID is also described as the composition of the PKT.NID and the PKT.LCID. The composition of these fields is described in Figure 8-21.




Figure 8-21 – PKT.CID structure

Table 8-4 – Packet header fields

Name	Length	Description
<i>Reserved</i>	3 bits	Always 0 for this version of the Recommendation. Reserved by ITU-T.
PKT.NAD	1 bit	No aggregation at destination <ul style="list-style-type: none"> If PKT.NAD=0 the packet may be aggregated with other packets at destination. If PKT.NAD=1 the packet may not be aggregated with other packets at destination.
PKT.PRIO	2 bits	Indicates a packet priority between 0 and 3
PKT.C	1 bits	Control <ul style="list-style-type: none"> If PKT.C=0 it is a data packet. If PKT.C=1 it is a control packet.
PKT.LCID / PKT.CTYPE	9 bits	Local connection identifier or control type <ul style="list-style-type: none"> If PKT.C=0, PKT.LCID represents the local connection identifier of a data packet. If PKT.C=1, PKT.CTYPE represents the type of the control packet.
PKT.SID	8 bits	Switch identifier <ul style="list-style-type: none"> If HDR.DO=0, PKT.SID represents the SID of the packet source. If HDR.DO=1, PKT.SID represents the SID of the packet destination.
PKT.LNID	14 bits	Local node identifier <ul style="list-style-type: none"> If HDR.DO=0, PKT.LNID represents the LNID of the packet source. If HDR.DO=1, PKT.LNID represents the LNID of the packet destination.
PKT.SPAD	1bit	Indicates if padding is inserted while encrypting payload. Note that this bit is only of relevance when security profile 1 (see clause 8.3.8.2.2) is used.
PKT.LEN	9 bits	Length of the packet payload in bytes

8.4.2.4 CRC

The CRC is the last field of the GPDU. It is 32 bits long. It is used to detect transmission errors. The CRC shall cover the concatenation of the SNA with the GPDU except for the CRC field itself.

The input polynomial $M(x)$ is formed as a polynomial whose coefficients are bits of the data being checked (the first bit to check is the highest order coefficient and the last bit to check is the coefficient of order zero). The generator polynomial for the CRC is $G(x)=x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^8+x^7+x^5+x^4+x^2+x+1$. The remainder $R(x)$ is calculated as the remainder from the division of $M(x) \cdot x^{32}$ by $G(x)$. The coefficients of the remainder will then be the resulting CRC.

8.4.3 Promotion needed PDU

If a node is disconnected and it does not have connectivity with any existing switch node, it shall send notifications to its neighbours to indicate the need for the promotion of any available terminal node. Figure 8-22 represents the promotion needed MAC PDU (PNPDU) that must be sent on an irregular basis in this situation.

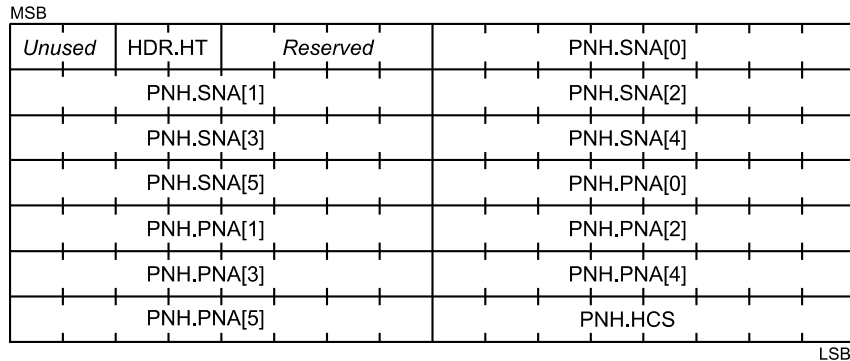


Figure 8-22 – Promotion need MAC PDU

Table 8-5 shows the promotion need MAC PDU fields.

Table 8-5 – Promotion need MAC PDU fields

Name	Length	Description
<i>Unused</i>	2 bits	Unused bits which are always 0; included for alignment with MAC_H field in the PDU header (see clause 7.4.3).
HDR.HT	2 bits	Header type HDR.HT = 1 for the promotion need MAC PDU
<i>Reserved</i>	4 bits	Always 0 for this version of the Recommendation. Reserved by ITU-T.
PNH.SNA	48 bits	Subnetwork address The EUI-48 of the base node of the subnetwork the service node is trying to connect to. FF:FF:FF:FF:FF:FF to ask for the promotion in any available subnetwork. SNA[0] is the most significant byte of the OUI/IAB and SNA[5] is the least significant byte of the extension identifier, as defined in: http://standards.ieee.org/regauth/oui/tutorials/EUI-48.html . The above notation is applicable to all EUI-48 fields in the specification.
PNH.PNA	48 bits	Promotion need address. The EUI-48 of the node that needs the promotion. It is the EUI-48 of the transmitter.
PNH.HCS	8 bits	Header check sequence. A field for detecting errors in the header. The transmitter shall calculate the PNH.HCS of the first 13 bytes of the header and insert the result into the PNH.HCS field (the last byte of the header). It shall be calculated as the remainder of the division (Modulo 2) of the polynomial $M(x) \cdot x^8$ by the generator polynomial $g(x) = x^8 + x^2 + x + 1$. $M(x)$ is the input polynomial, which is formed by the bit sequence of the header excluding the PNH.HCS field, and the MSB of the bit sequence is the coefficient of the highest order of $M(x)$.

As it is always transmitted by unsynchronized nodes and therefore, prone to creating collisions, it is a special reduced size header. It is broadcast to any other terminal node and shall therefore be transmitted with the most robust scheme of the PHY layer.

8.4.4 Beacon PDU

Beacon PDU (BPDU) is transmitted by every switch device on the subnetwork, including the base node. The purpose of this PDU is to circulate information on MAC frame structure and therefore channel access to all devices that are part of this subnetwork. The BPDU is transmitted at definite fixed intervals of time and is also used as a synchronization mechanism by service nodes. Figure 8-23 below shows contents of a beacon transmitted by the base node and each switch device.

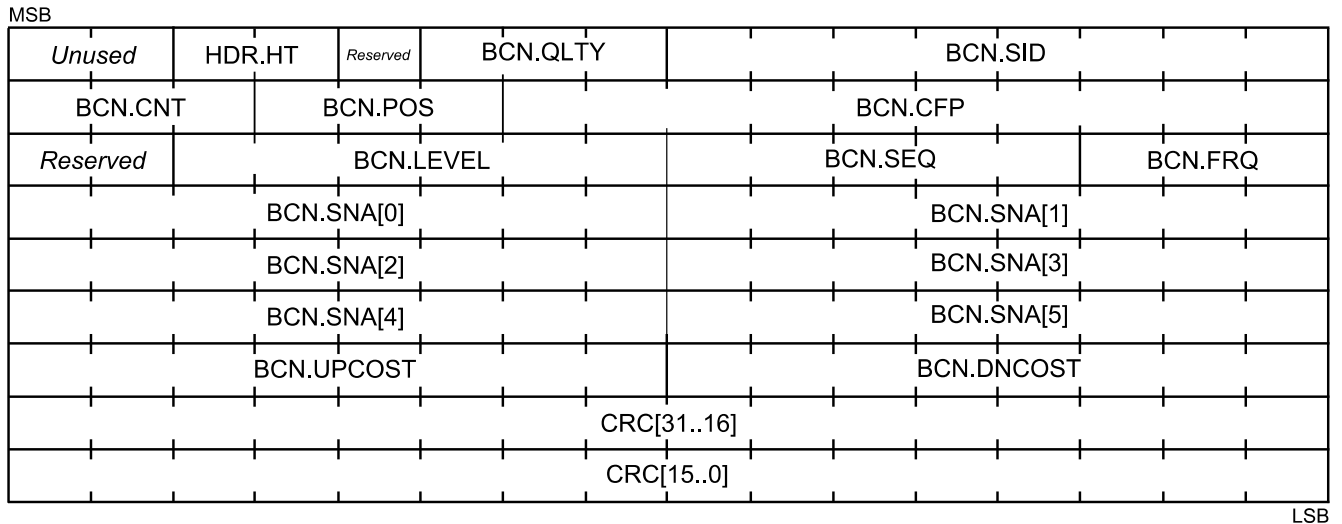


Figure 8-23 – Beacon PDU structure

Table 8-6 shows the beacon PDU fields.

Table 8-6 – Beacon PDU fields

Name	Length	Description
Unused	2 bits	Unused bits which are always 0; included for alignment with MAC_H field in the PDU header (see clause 7.4.3).
HDR.HT	2 bits	Header type HDR.HT = 2 for Beacon PDU
Reserved	1 bit	Always 0 for this version of the Recommendation. Reserved by ITU-T.
BCN.QLTY	3 bits	Quality of round-trip connectivity from this switch node to the base node. BCN.QLTY=7 for best quality (base node or very good switch node), BCN.QLTY=0 for worst quality (switch having unstable connection to subnetwork)
BCN.SID	8 bits	Switch identifier of transmitting switch
BCN.CNT	3 bits	Number of beacon slots in this frame
BCN.SLT	3 bits	Beacon slot in which this BPDU is transmitted BCN.SLT=0 is reserved for the base node
BCN.CFP	10 bits	Offset of CFP from start of frame BCN.CFP=0 indicates absence of CFP in a frame.
Reserved	1 bit	Always 0 for this version of the Recommendation. Reserved by ITU-T.
BCN.LEVEL	6 bits	Hierarchy of transmitting switch in the subnetwork

Table 8-6 – Beacon PDU fields

Name	Length	Description
BCN.SEQ	5 bits	Sequence number of this BPDU in the superframe. Incremented for every beacon the base node sends and is propagated by switch through its BPDU so that the entire subnetwork has the same notion of sequence number at a given time.
BCN.FRQ	3 bits	Transmission frequency of this BPDU. Values are interpreted as follows: 0 = 1 beacon every frame 1 = 1 beacon every 2 frames 2 = 1 beacon every 4 frames 3 = 1 beacon every 8 frames 4 = 1 beacon every 16 frames 5 = 1 beacon every 32 frames 6 = Reserved by ITU-T 7 = Reserved.by ITU-T.
BCN.SNA	48 bits	Subnetwork identifier in which the switch transmitting this BPDU is located.
BCN.UPCOST	8 bits	Total uplink cost from the transmitting switch node to the base node. The cost of a single hop is calculated based on the modulation scheme used on that hop in the uplink direction. Values are derived as follows: 8PSK = 0 QPSK = 1 BPSK = 2 8PSK_F = 1 QPSK_F = 2 BPSK_F = 4. The base node will transmit in its beacon a BCN.UPCOST of 0. A switch node will transmit in its beacon the value of BCN.UPCOST received from its upstream switch node, plus the cost of the upstream uplink hop to its upstream switch. When this value is larger than what can be held in BCN.UPCOST the maximum value of BCN.UPCOST should be used.
BCN.DNCOST	8 bits	Total downlink cost from the base node to the transmitting switch node. The cost of a single hop is calculated based on the modulation scheme used on that hop in the downlink direction. Values are derived as follows: 8PSK 0 QPSK 1 BPSK 2 8PSK_F 1 QPSK_F 2 BPSK_F 4. The base node will transmit in its beacon a BCN.DNCOST of 0. A switch node will transmit in its beacon the value of BCN.DNCOST received from its upstream switch node, plus the cost of the upstream downlink hop from its upstream switch. When this value is larger than what can be held in BCN.DNCOST the maximum value of BCN.DNCOST should be used.
CRC	32 bits	The CRC shall be calculated with the same algorithm as the one defined for the CRC field of the MAC PDU (see clause 8.4.2.4 for details). This CRC shall be calculated over the complete BPDU except for the CRC field itself.

The BPDUs are also used to detect when the uplink switch is no longer available either by a change in the characteristics of the medium or because of failure etc. If a service node fails to receive $N_{\text{miss-beacon}}$ in a row it should declare the link to its switch as unusable. The service node should stop sending beacons itself if it is acting as a switch. It should close all existing MAC connections. The service node then enters the initial unregistered functional state and searches for a subnetwork join. This mechanism complements the keep-alive mechanism which is used by a base node and its switches to determine when a service node is lost.

8.4.5 MAC control packets

8.4.5.1 General

MAC control packets enable a service node to communicate control information with their switch node, base node and vice versa. A control packet is transmitted as a GPDU and is identified with the PKT.C bit set to 1 (See clause 8.4.2 for more information about the fields of the packets).

There are several types of control messages. Each control message type is identified by the field PKT.CTYPE. Table 8-7 lists the types of control messages. The packet payload (see clause 8.4.2.3) shall contain the information carried by the control packets. This information differs depending on the packet type.

Table 8-7 – MAC control packet types

Type (PKT.CTYPE)	Packet name	Packet description
1	REG	Registration management
2	CON	Connection management
3	PRO	Promotion management
4	BSI	Beacon slot indication
5	FRA	Frame structure change
6	CFP	Contention-free period request
7	ALV	Keep-alive
8	MUL	Multicast management
9	PRM	PHY robustness management
10	SEC	Security information

8.4.5.2 Control packet retransmission

For recovery from lost control messages, a retransmit scheme is defined. MAC control transactions comprising of exchange of more than one control packet may follow the retransmission mechanism described in this clause.

The retransmission scheme shall be applied to the following packets when they require a response:

- CON_REQ_S, CON_REQ_B
- CON_CLS_S, CON_CLS_B
- REG_RSP
- PRO_REQ_B
- BSI_IND
- MUL_JOIN_S, MUL_JOIN_B
- MUL_LEAVE_S, MUL_LEAVE_B.

Devices involved in a MAC control transaction using the retransmission mechanism shall keep count of the number of times a message has been retransmitted and maintain a retransmit timer.

At the requester of a control message transaction:

- When the first message in a transaction is transmitted, the retransmit timer is started with the value *macCtlReTxTimer* and the retransmit count is set to 0.

If a response message is received the retransmit timer is stopped and the transaction is considered complete. Note that it is possible to receive further response messages. These would be messages that encountered network delays.

- If the retransmit timer expires, the retransmit counter is incremented. If the retransmit counter is less than *macMaxCtlReTx* the control message is retransmitted. If the counter is equal to the maximum number of retransmits, the failure result corresponding to the respective MAC-SAP should be returned to the calling entity. Implementations may also choose to inform their local management entity of such a failure. If the retransmission is done by the service node, the device should return to the disconnected state.

At the responder of a control message transaction:

- The receiver of a message must determine itself if this message is a retransmit. If so, no local action is needed other than sending a reply to the response.

If the received message is not a retransmit, the message should be processed and a response returned to the sender.

- For transactions which use three messages in the transaction, e.g., promotion as shown in clause 8.6.3, the responder should perform retransmits in exactly the same way as the requester. This ensures that if the third message in the transaction is lost, the message will be retried and the transaction completed.

The following message sequence charts show some examples of retransmission. Figure 8-24 shows two successful transactions without requiring retransmits.

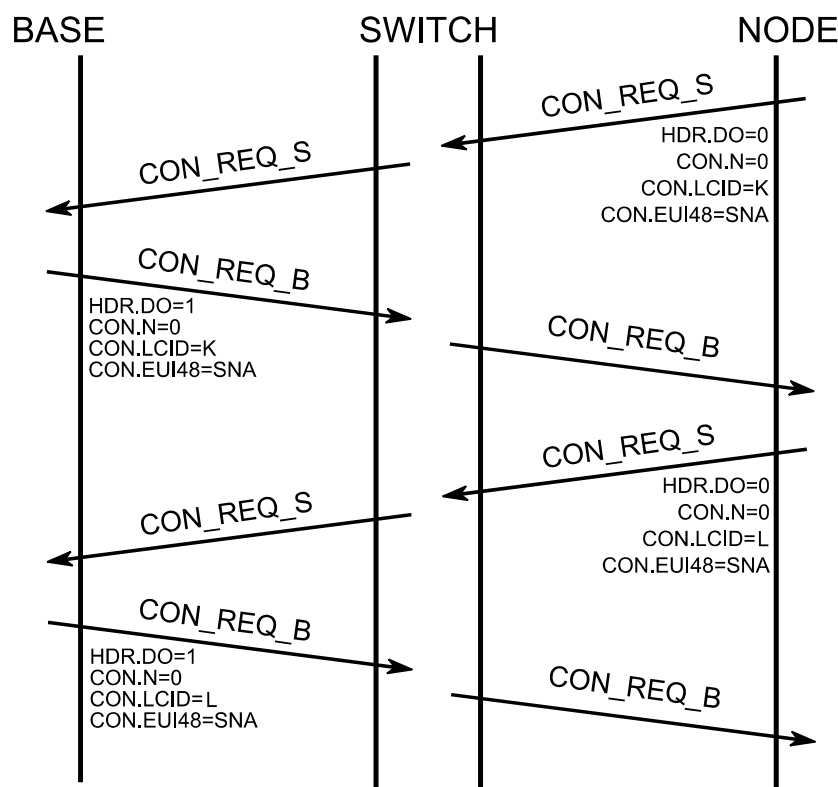


Figure 8-24 – Two transactions without requiring retransmits

Figure 8-25 shows a more complex example, where messages are lost in both directions causing multiple retransmits before the transaction completes.

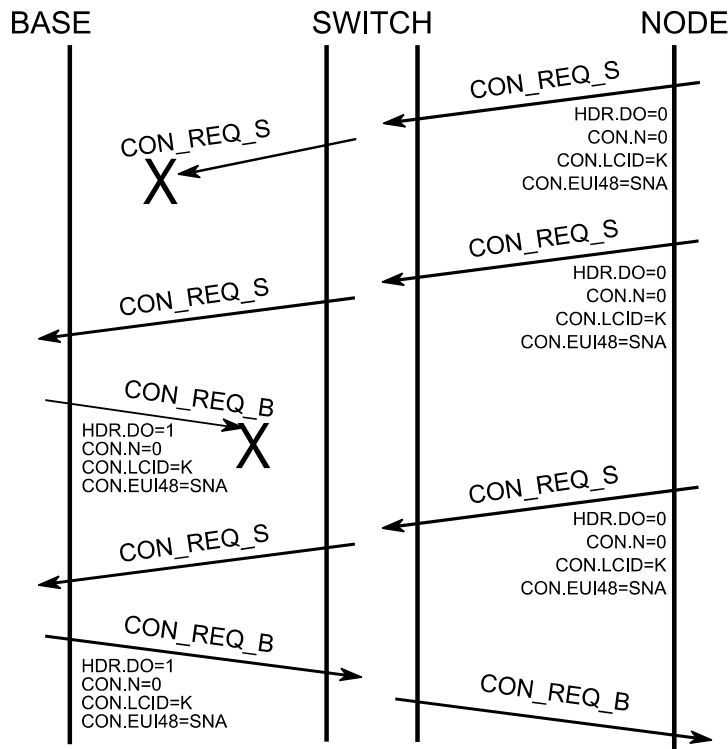


Figure 8-25 – Transaction with packet loss requiring retransmits

Figure 8-26 shows the case of a delayed response causing duplication at the initiator of the control transaction.




Figure 8-26 – Duplicate packet detection and elimination

8.4.5.3 REG control packet (PKT.CTYPE=1)

This control packet is used to negotiate the registration process. The description of data fields of this control packet is described in Table 8-8 and Figure 8-27. The meaning of the packets differs depending on the direction of the packet. This packet interpretation is explained in Table 8-9. These packets are used during the registration and unregistration processes, as explained in clauses 8.6.1 and 8.6.2.

The PKT.SID field is used in this control packet as the switch where the service node is registering. The PKT.LNID field is used in this control packet as the local node identifier being assigned to the service node during the registration process negotiation.

The REG.CAP_PA field is used to indicate the packet aggregation capability as discussed in clause 8.3.7. In the uplink direction, this field is an indication from the registering terminal node about its own capabilities. For the downlink response, the base node evaluates whether or not all the devices in the cascaded chain from itself to this terminal node have the packet-aggregation capability. If they do, the base node shall set REG.CAP_PA=1; otherwise REG.CAP_PA=0.

MSB

REG.N	REG.R	REG.SPC	Reserved	REG.CAP_SW	REG.CAP_PA	REG.CAP_CFP	REG.CAP_DC	REG.CAP_MC	REG.CAP_PRM	REG.CAP_ARQ	REG.TIME
REG.EUI48[47..40]						REG.EUI48[39..32]					
REG.EUI48[31..24]						REG.EUI48[23..16]					
REG.EUI48[15..8]						REG.EUI48[7..0]					
REG.SNK[127..112]											
REG.SNK[111..96]											
REG.SNK[95..80]											
REG.SNK[79..64]											
REG.SNK[63..48]											
REG.SNK[47..32]											
REG.SNK[31..16]											
REG.SNK[15..0]											
REG.AUK[127..112]											
REG.AUK[111..96]											
REG.AUK[95..80]											
REG.AUK[79..64]											
REG.AUK[63..48]											
REG.AUK[47..32]											
REG.AUK[31..16]											
REG.AUK[15..0]											

G.9904(12)_F8-27

Figure 8-27 – REG control packet structure

Table 8-8 – REG control packet fields

Name	Length	Description
REG.N	1 bit	Negative <ul style="list-style-type: none">• REG.N=1 for the negative register• REG.N=0 for the positive register (see Table 8-9)
REG.R	1 bit	Roaming <ul style="list-style-type: none">• REG.R=1 if node already registered and wants to perform roaming to another switch.• REG.R=0 if node not yet registered and wants to perform a clear registration process.
REG.SPC	2 bits	Security profile capability for data PDUs: <ul style="list-style-type: none">• REG.SPC=0 No encryption capability• REG.SPC=1 Security profile 1 capable device• REG.SPC=2 Security profile 2 capable device (not yet specified)• REG.SPC=3 Security profile 3 capable device (not yet specified).
<i>Reserved</i>	2 bits	Reserved by ITU-T. Should be set to 0 for this version of the protocol.
REG.CAP_SW	1 bit	Switch capable 1 if the device is able to behave as a switch node 0 if the device is not
REG.CAP_PA	1 bit	Packet aggregation capability 1 if the device has packet aggregation capability (uplink) if the data transit path to the device has packet aggregation capability (downlink); 0 otherwise.
REG.CAP_CFP	1 bit	Contention-free period capability 1 if the device is able to perform the negotiation of the CFP; 0 if the device cannot use the contention-free period in a negotiated way.
REG.CAP_DC	1 bit	Direct connection capability 1 if the device is able to perform direct connections 0 if the device is not able to perform direct connections
REG.CAP_MC	1 bit	Multicast capability 1 if the device is able to use multicast for its own communications 0 if the device is not able to use multicast for its own communications
REG.CAP_PRM	1 bit	PHY robustness management capable 1 if the device is able to perform PHY robustness management 0 if the device is not able to perform PHY robustness management
REG.CAP_ARQ	1 bit	ARQ capable 1 if the device is able to establish ARQ connections 0 if the device is not able to establish ARQ connections

Table 8-8 – REG control packet fields

Name	Length	Description
REG.TIME	3 bits	Time to wait for an ALV_B messages before assuming the service node has been unregistered by the base node. For all messages except REG_RSP this field should be set to 0. For REG_RSP its value means: ALV.TIME = 0 => 32 seconds ALV.TIME = 1 => 64 seconds ALV.TIME = 2 => 128 seconds ~ 2.1 minutes ALV.TIME = 3 => 256 seconds ~ 4.2 minutes ALV.TIME = 4 => 512 seconds ~ 8.5 minutes ALV.TIME = 5 => 1024 seconds ~ 17.1 minutes ALV.TIME = 6 => 2048 seconds ~ 34.1 minutes ALV.TIME = 7 => 4096 seconds ~ 68.3 minutes.
REG.EUI-48	48 bit	EUI-48 of the node EUI-48 of the node requesting the registration
REG.SNK	128 bits	Encrypted subnetwork key that shall be used to derive the subnetwork working key
REG.AUK	128 bits	Encrypted authentication key. This is a random sequence meant to act as an authentication mechanism.

Table 8-9 – REG control packet types

Name	HDR.DO	PKT.LNID	REG.N	REG.R	Description
REG_REQ	0	0x3FFF	0	R	Registration request <ul style="list-style-type: none"> If R=0 any previous connection from this node should be lost; If R=1 any previous connection from this node should be maintained.
REG_RSP	1	< 0x3FFF	0	R	Registration response. This packet assigns the PCK.LNID to the service node.
REG_ACK	0	< 0x3FFF	0	R	Registration acknowledged by the service node.
REG_REJ	1	0x3FFF	1	0	Registration rejected by the base node.
REG_UNR_S	0	< 0x3FFF	1	0	<ul style="list-style-type: none"> After a REG_UNR_B: unregistration acknowledge Alone: unregistration request initiated by the node
REG_UNR_B	1	< 0x3FFF	1	0	<ul style="list-style-type: none"> After a REG_UNR_S: unregistration acknowledge Alone: unregistration request initiated by the base node

Fields REG.SNK and REG.AUK are of significance only for REG_RSP and REG_ACK messages with security profile 1 (REG.SCP=1). For all other message-exchange variants using the REG control packet, these fields shall not be present reducing the length of payload.

In REG_RSP message, the REG.SNK and REG.AUK shall always be inserted encrypted with WK0.

In the REG_ACK message, the REG.SNK field shall be set to zero. The contents of the REG.AUK field shall be derived by decrypting the received REG_RSP message with WK0 and re-encrypting the decrypted REG.AUK field with SWK derived from the decrypted REG.SNK and random sequence previously received in SEC control packets.

8.4.5.4 CON control packet (PKT.CTYPE = 2)

This control packet is used for negotiating the connections. The description of the fields of this packet is given in Table 8-10 and Figure 8-28. The meaning of the packet differs depending on the direction of the packet and on the values of the different types. Table 8-11 shows the different interpretation of the packets. The packets are used during the connection establishment and closing. These processes are explained in more detail in clause 8.6.6.

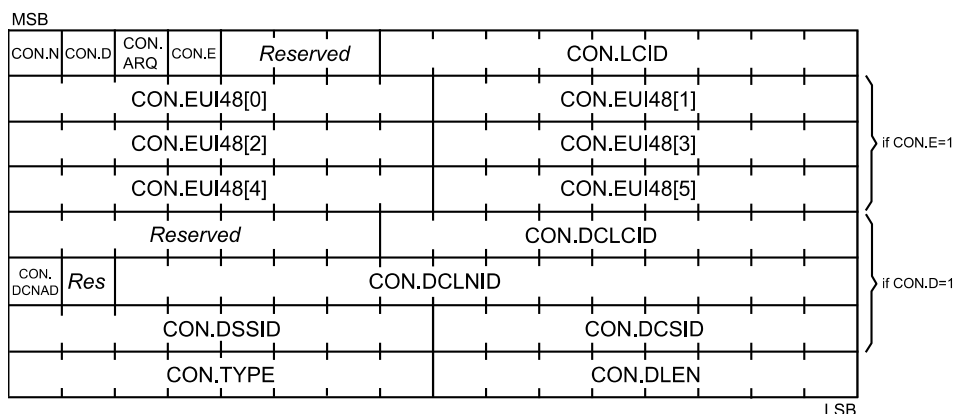


Figure 8-28 – CON control packet structure

Note that Figure 8-28 shows the complete message with all optional parts. When CON.D is 0, CON.DCNAD, CON.DSSID, CON.DCLNID, CON.DCLID, CON.DCSID and the reserved field between CON.DCNAD and CON.DSSID will not be present in the message. Thus, the message will be 6 octets smaller. Similarly, when CON.E is zero, the field CON.EUI-48 will not be present, making the message 6 octets smaller.

Table 8-10 – CON control packet fields

Name	Length	Description
CON.N	1 bit	Negative <ul style="list-style-type: none"> • CON.N=1 for the negative connection • CON.N=0 for the positive connection
CON.D	1 bit	Direct connection <ul style="list-style-type: none"> • CON.D=1 if information about direct connection is carried by this packet • CON.D=0 if information about direct connection is not carried by this packet
CON.ARQ	1 bit	ARQ mechanism enable <ul style="list-style-type: none"> • CON.ARQ=1 if ARQ mechanism is enabled for this connection • CON.ARQ=0 if ARQ mechanism is not enabled for this connection

Table 8-10 – CON control packet fields

Name	Length	Description
CON.E	1 bit	EUI-48 presence <ul style="list-style-type: none"> • CON.E = 1 to have a CON.EUI-48 • CON.E = 0 to not have a CON.EUI-48 so that this connection establishment is for reaching the base node CL.
Reserved	3 bits	Reserved by ITU-T. This shall be 0 for this version of the protocol.
CON.LCID	9 bits	Local connection identifier The LCID is reserved in the connection request. LCIDs from 0 to 255 are assigned by the connection requests initiated by the base node. LCIDs from 256 to 511 are assigned by the connection requests initiated by the local node. This is the identifier of the connection being managed with this packet. This is not the same as the PKT.LCID of the generic header, which does not exist for control packets.
CON.EUI-48	48 bits (Present if CON.E=1)	EUI-48 of destination/source service node/base node for connection request When not performing a directed connection, this field should not be included. When performing a directed connection, it may contain the SNA, indicating that the base node convergence layer should determine the EUI-48. <ul style="list-style-type: none"> • CON.D = 0, Destination EUI-48 • CON.D = 1, Source EUI-48
Reserved	7 bits (Present if CON.D=1)	Reserved by ITU-T. This shall be 0 for this version of the protocol.
CON.DCLCID	9 bits (Present if CON.D=1)	Direct connection LCID This field represents the LCID of the connection identifier to which the one being established shall be directly switched.
CON.DCNAD	1 bit (Present if CON.D=1)	Reserved by ITU-T. Direct connection not aggregated at destination. This field represents the content of the PKT.NAD field after a direct connection switch operation.
Reserved	1 bits (Present if CON.D=1)	Reserved by ITU-T. This shall be 0 for this version of the protocol.
CON.DCLNID	14 bits (Present if CON.D=1)	Direct connection LNID This field represents the LNID part of the connection identifier to which the one being established shall be directly switched.
CON.DSSID	8 bits (Present if CON.D=1)	Direct switch SID This field represents the SID of the switch that should learn this direct connection and perform direct switching.
CON.DCSID	8 bits (Present if CON.D=1)	Direct connection SID This field represents the SID part of the connection identifier to which the one being established shall be directly switched.

Table 8-10 – CON control packet fields

Name	Length	Description
CON.TYPE	8 bits	Connection type The connection type (see Annex C) specifies the convergence layer to be used for this connection. They are treated transparently through the MAC common part sublayer and are used only to identify which convergence layer may be used.
CON.DLEN	8 bits	Length of CON.DATA field in bytes
CON.DATA	(variable)	Connection specific parameters These connection specific parameters are convergence layer specific. They should be defined in each convergence layer to define the parameters that are specific to the connection. These parameters are handled in a transparent way by the common part sublayer.

Table 8-11 – CON control packet types

Name	HDR.DO	CON.N	Description
CON_REQ_S	0	0	Connection establishment request initiated by the service node.
CON_REQ_B	1	0	The base node will consider that the connection is established with the identifier CON.LCID. <ul style="list-style-type: none"> • After a CON_REQ_S: connection accepted • Alone: Connection establishment request
CON_CLS_S	0	1	The service node considers this connection closed: <ul style="list-style-type: none"> • After a CON_REQ_B: connection rejected by the node • After a CON_CLS_B: connection closing acknowledge • Alone: connection closing request
CON_CLS_B	1	1	The base node will consider that the connection is no longer established. <ul style="list-style-type: none"> • After a CON_REQ_S: connection establishment rejected by the base node • After a CON_CLS_S: connection closing acknowledge • Alone: connection closing request

8.4.5.5 PRO control packet (PKT.CTYPE = 3)

This control packet is used to promote a service node from the terminal function to switch function. The description of the fields of this packet is given in Table 8-12, Figures 8-29 and 8-30. The meaning of the packet differs depending on the direction of the packet and on the values of the different types. Table 8-13 shows the different interpretation of the packets. The promotion process is explained in more detail in 8.6.3.

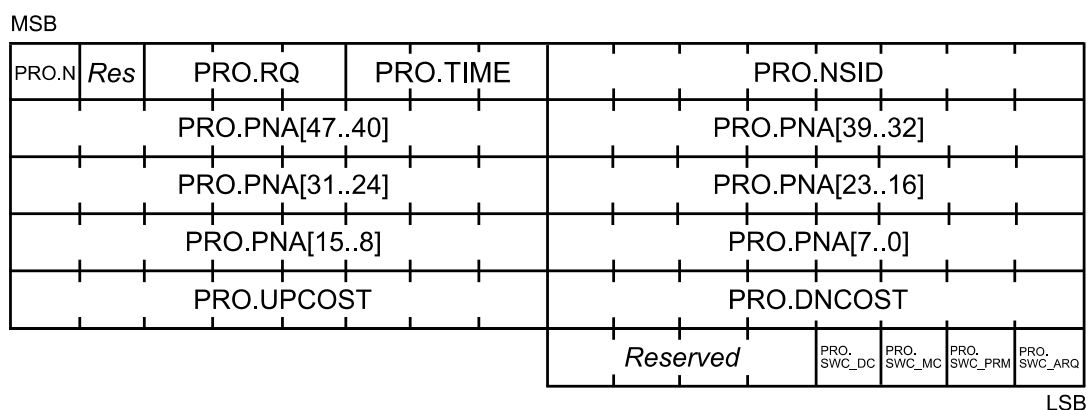


Figure 8-29 – PRO_REQ_S control packet structure




Figure 8-30 – PRO control packet structure

Note that Figure 8-29 includes all fields as used by a PRO_REQ_S message. All other messages are much smaller, containing only PRO.N, PRO.RC, PRO.TIME and PRO.NSID as shown in Figure 8-30.

Table 8-12 – PRO control packet fields

Name	Length	Description
PRO.N	1 bit	Negative PRO.N=1 for the negative promotion PRO.N=0 for the positive promotion
Reserved	1 bit	Reserved by ITU-T. This shall be 0 for this version of the protocol.
PRO.RQ	3 bits	Receive quality of the PNPDU message received from the service node requesting the terminal to promote.
PRO.TIME	3 bits	The ALV.TIME which is being used by the terminal which will become a switch. Upon receipt of this time in a PRO_REQ_B the service node should reset the keep-alive timer in the same way as receiving an ALV_B.
PRO.NSID	8 bits	New switch identifier. This is the assigned switch identifier of the node whose promotion is being managed with this packet. This is not the same as the PKT.SID of the packet header, which must be the SID of the switch this node is connected to, as a terminal node.
PRO.PNA	0 or 48 bits	Promotion need address, contains the EUI-48 of the terminal requesting the service node promotes to become a switch. This field is only included in the PRO_REQ_S message.
PRO.UPCOST	0 or 8 bits	Total uplink cost from the terminal node to the base node. This value is calculated in the same way a switch node calculates the value it places into its own beacon PDU. This field is only included in the PRO_REQ_S message.

Table 8-12 – PRO control packet fields

Name	Length	Description
PRO.DNCOST	0 or 8 bits	Total downlink cost from the base node to the terminal node. This value is calculated in the same way a switch node calculates the value it places into its own beacon PDU. This field is only included in the PRO_REQ_S message.
Reserved	4 bits	Reserved by ITU-T. Should be set to 0 for this version of the protocol.
PRO.SWC_DC	1 bit	Direct connection switching capability 1 if the device is able to behave as a direct switch in direct connections 0 otherwise
PRO.SWC_MC	1 bit	Multicast switching capability 1 if the device is able to manage the multicast traffic when behaving as a switch 0 otherwise
PRO.SWC_PRM	1 bit	PHY robustness management switching capability 1 if the device is able to perform PRM for the terminal nodes when behaving as a switch. 0 if the device is not able to perform PRM when behaving as a switch.
PRO.SWC_ARQ	1 bit	ARQ buffering switching capability 1 if the device is able to perform buffering for ARQ connections while switching. 0 if the device is not able to perform buffering for ARQ connections while switching.

Table 8-13 – PRO control packet types

Name	HDR.DO	PRO.N	PRO.NSID	Description
PRO_REQ_S	0	0	0xFF	Promotion request initiated by the service node.
PRO_REQ_B	1	0	< 0xFF	The base node will consider that the service node has promoted with the identifier PRO.NSID. <ul style="list-style-type: none"> • After a PRO_REQ: promotion accepted • Alone: promotion request initiated by the base node
PRO_ACK	0	0	< 0xFF	Promotion acknowledge
PRO_REJ	1	1	0xFF	The base node will consider that the service node is demoted. It is sent after a PRO_REQ to reject it.
PRO_DEM_S	0	1	< 0xFF	The service node considers that it is demoted: <ul style="list-style-type: none"> • After a PRO_DEM_B: demotion accepted • After a PRO_REQ_B: promotion rejected • Alone: demotion request.
PRO_DEM_B	1	1	< 0xFF	The base node considers that the service node is demoted. <ul style="list-style-type: none"> • After a PRO_DEM_S: demotion accepted • Alone: demotion request

BSI control packet (PKT.CTYPE = 4)

The beacon slot information (BSI) control packet is only used by the base node and switch nodes. It is used to exchange information that is further used by a switch node to transmit its beacon. The description of the fields of this packet is given in Table 8-14 and Figure 8-31. The meaning of the packet differs depending on the direction of the packet and on the values of the different types. Table 8-15 represents the different interpretation of the packets. The promotion process is explained in more detail in clause 8.6.3.




Figure 8-31 – BSI control packet structure

Table 8-14 – BSI control packet fields

Name	Length	Description
Reserved	5 bits	Reserved by ITU-T. In this version, this field should be initialized to 0.
BSI.FRQ	3 bits	Transmission frequency of Beacon slot, encoded as: FRQ = 0 => 1 beacon every frame FRQ = 1 => 1 beacon every 2 frames FRQ = 2 => 1 beacon every 4 frames FRQ = 3 => 1 beacon every 8 frames FRQ = 4 => 1 beacon every 16 frames FRQ = 5 => 1 beacon every 32 frames FRQ = 6 => Reserved by ITU-T FRQ = 7 => Reserved by ITU-T.
BSI.SLT	3 bits	Beacon slot to be used by target switch
BSI.SEQ	5 bits	The Beacon sequence number when the specified change takes effect.

Table 8-15 – BSI control message types

Name	HDR.DO	Description
BSI_ACK	0	Acknowledgement of receipt of BSI control message
BSI_IND	1	Beacon-slot change command

8.4.5.6 FRA control packet (PKT.CTYPE = 5)

This control packet is broadcast from the base node and relayed by all switch nodes to the entire subnetwork. It is used to circulate information on the change of frame structure at a specific time in future. The description of fields of this packet is given in Table 8-16 and Figure 8-32. Table 8-17 shows the different interpretation of the packets.




Figure 8-32 – FRA control packet structure

Table 8-16 – FRA control packet fields

Name	Length	Description
FRA.TYP	2 bits	0: Beacon count change 1: CFP duration change
Reserved	4 bits	Reserved by ITU-T. In this version, this field should be initialized to 0.
FRA.CFP	10 bits	Offset of CFP from the start of frame
FRA.SEQ	5 bits	The beacon sequence number when the specified change takes effect.
FRA.BCN	3 bits	Number of beacons in a frame

Table 8-17 – FRA control packet types

Name	FRA.TYP	Description
FRA_BCN_IND	0	Indicates changes to frame structure due to change in the beacon-slot count
FRA_CFP_IND	1	Indicates changes to frame structure due to change in CFP duration as a result of grant of CFP or end of CFP period for any requesting service node in the subnetwork.

8.4.5.7 CFP control packet (PKT.CTYPE = 6)

This control packet is used for dedicated contention-free channel access time allocation to individual terminal or switch nodes. The description of the fields of this packet is given in Table 8-18 and Figure 8-33. The meaning of the packet differs depending on the direction of the packet and on the values of the different types. Table 8-19 represents the different interpretation of the packets.




Figure 8-33 – CFP control packet structure

Table 8-18 – CFP control message fields

Name	Length	Description
CFP.N	1 bit	0: denial of allocation/de-allocation request 1: acceptance of allocation/de-allocation request
CFP.DIR	1 bit	Indicate direction of allocation. 0: allocation is applicable to uplink (towards base node) direction 1: allocation is applicable to downlink (towards service node) direction
CFP.SEQ	5 bits	The beacon sequence number when the specified change takes effect.
CFP.LCID	9 bits	LCID of requesting connection
CFP.LEN	7 bits	Length (in symbols) of requested/allocated channel time per frame.
CFP.POS	9 bits	Offset (in symbols) of allocated time from beginning of frame.

Table 8-18 – CFP control message fields

Name	Length	Description
CFP.TYPE	2 bits	0: channel allocation packet 1: channel de-allocation packet 2: channel change packet
CFP.LNID	14 bits	LNID of service node that is the intended user of the allocation.

Table 8-19 – CFP control packet types

Name	CFP. TYP	HDR. DO	Description
CFP_ALC_REQ_S	0	0	Service node makes channel allocation request
CFP_ALC_IND	0	1	<ul style="list-style-type: none"> After a CFP_ALC_REQ_S: requested channel is allocated Alone: unsolicited channel allocation by base node
CFP_ALC_REJ	0	1	Requested channel allocation is denied
CFP_DALC_REQ	1	0	Service node makes channel de-allocation request
CFP_DALC_RSP	1	1	Base node confirms de-allocation
CFP_CHG_IND	2	1	Change of location of allocated channel within the CFP

8.4.5.8 ALV control packet (PKT.CTYPE = 7)

The ALV control message is used for keep-alive signalling between a service node, the service nodes above it and the base node. The message exchange is bidirectional, that is, a message is periodically exchanged in each direction. The structure of these messages are shown in Figure 8-34 and Table 8-20. The different keep-alive message types are shown in Table 8-21. These messages are sent periodically, as described in clause 8.6.5.




Figure 8-34 – ALV control packet structure

Table 8-20 – ALV control message fields

Name	Length	Description
ALV.RXCNT	3 bits	Modulo 8 counter to indicate the number of received ALV messages.
ALV.TXCNT	3 bits	Modulo 8 counter to indicate the number of transmitted ALV messages.
Reserved	7 bits	Reserved by ITU-T. Should always be encoded as 0 in this version of the Recommendation.
ALV.TIME	3 bits	Time to wait for ALV_B messages before assuming the service node has been unregistered by the base node. ALV.TIME = 0 => 32 seconds ALV.TIME = 1 => 64 seconds ALV.TIME = 2 => 128 seconds ~ 2.1 minutes ALV.TIME = 3 => 256 seconds ~ 4.2 minutes ALV.TIME = 4 => 512 seconds ~ 8.5 minutes ALV.TIME = 5 => 1024 seconds ~ 17.1 minutes ALV.TIME = 6 => 2048 seconds ~ 34.1 minutes ALV.TIME = 7 => 4096 seconds ~ 68.3 minutes
ALV.SSID	8 bits	For a terminal, this should be 0xFF. For a switch, this is its switch identifier.

Table 8-21 – Keep-alive control packet types

Name	HDR.DO	Description
ALV_S	0	Keep-alive message from a service node
ALV_B	1	Keep-alive message from the base node

8.4.5.9 MUL control packet (PKT.CTYPE = 8)

The MUL message is used to control multicast group membership. The structure of this message and the meanings of the fields are described in Table 8-22 and Figure 8-35. The message can be used in different ways as described in Table 8-23.




Figure 8-35 – MUL control packet structure

Table 8-22 – MUL control message fields

Name	Length	Description
MUL.N	1 bit	Negative <ul style="list-style-type: none"> • MUL.N = 1 for the negative multicast connection, i.e., multicast group leave. • MUL.N = 0 for the positive multicast connection, i.e., multicast group join.
Reserved	6 bits	Reserved by ITU-T. This shall be 0 for this version of the protocol.
MUL.LCID	9 bits	Local connection identifier The LCID indicates which multicast distribution group is being managed with this message.
MUL.TYPE	8 bits	Connection type The connection type specifies the convergence layer to be used for this connection. They are treated transparently through the MAC common part sublayer and are used only to identify which convergence layer may be used. See Annex C.
MUL.DLEN	8 bits	Length of data in bytes in the MUL.DATA field
MUL.DATA	(variable)	Connection specific parameters These connection specific parameters are convergence layer specific. They should be defined in each convergence layer to define the parameters that are specific to the connection. These parameters are handled in a transparent way by the common part sublayer.

Table 8-23 – MUL control message types

Name	HDR.DO	MUL.N	Description
MUL_JOIN_S	0	0	Multicast group join request initiated by the service node, or an acknowledgement when sent in response to a MUL_JOIN_B.
MUL_JOIN_B	1	0	The base node will consider that the group has been joined with the identifier MUL.LCID. <ul style="list-style-type: none"> • After a MUL_JOIN_S: join accepted • Alone: group join request
MUL_LEAVE_S	0	1	The service node leaves the multicast group: <ul style="list-style-type: none"> • After a MUL_JOIN_B: join rejected by the node • After a MUL_LEAVE_B: group leave acknowledge • Alone: group leave request.
MUL_LEAVE_B	1	1	The base node will consider that the service node is no longer a member of the multicast group. <ul style="list-style-type: none"> • After a MUL_JOIN_S: group join rejected by the base node • After a MUL_LEAVE_S: group leave acknowledge • Alone: group leave request.

8.4.5.10 PRM control packet (PKT.CTYPE = 9)

The PHY robustness management packets are used to control the parameters that affect the robustness and efficiency of the PHY. These packets are sent to notify to the peer of the need to improve robustness of the transmission, or to notify the peer that the reception quality is so good that a less robust and so more efficient modulation scheme can be transmitted.

The fields of the PRM control packet are described in Table 8-24 and Figure 8-36 and the types of messages are described in Table 8-25.




Figure 8-36 – PHY control packet structure

Table 8-24 – PRM control message fields

Name	Length	Description
PRM.R	1 bit	Response <ul style="list-style-type: none"> • PRM.R=1 if this message is a response • PRM.R=0 if this message is a request
PRM.N	1 bit	Negative <ul style="list-style-type: none"> • PRM.N=1 if the operation could not be performed • PRM.N=0 if the operation was performed
Reserved	3 bits	Reserved by ITU-T. Should always be encoded as 0 in this version of the Recommendation.
PRM.SNR	3 bits	Indicates the SNR at the end that initiates a change request, obtained using PHY_SNR primitive (see clause 7.10.3.10).

Table 8-25 – PRM control message types

Name	PRM.R	PRM.N	Description
PRM_REQ	0	0	PHY modulation management request
PRM_ACK	1	0	PHY modulation management acknowledge
PRM_REJ	1	1	PHY modulation management rejected

8.4.5.11 SEC control packet (PKT.CTYPE = 10)

The SEC control message is broadcast unencrypted by the base node and all switch nodes to the rest of the subnetwork in order to circulate the random sequence used to generate working keys. The random sequence used by devices in a subnetwork is dynamic and changes from time to time to ensure a robust security framework. The structure of this message is shown in Table 8-26 and Figure 8-37. Further details of security mechanisms are given in clause 8.3.8.




Figure 8-37 – SEC control packet structure

Table 8-26 – SEC control message fields

Name	Length	Description
SEC.RAN	128 bits	Random sequence to be used to derive WK.
SEC.SNK	128 bits	Random sequence to be used to derive SWK.
Reserved	2 bits	Reserved by ITU-T. Should always be encoded as 0 in this version of the Recommendation.
SEC.USE	1 bits	When 1 indicates the random sequences are already in use. When 0 indicates that SE.SEQ should be used to determine when to start using these random sequences.
SEC.SEQ	5 bits	The beacon sequence number when the specified change takes effect.

8.5 MAC service access point

8.5.1 General

The MAC service access point provides several primitives to allow the convergence layer to interact with the MAC layer. This clause aims to explain how the MAC may be used. An implementation of the MAC may not use all the primitives listed here; it may use other primitives; or it may have a function-call based interface rather than message-passing, etc. These are all implementation issues which are beyond the scope of this Recommendation.

The .request primitives are passed from the CL to the MAC to request the initiation of a service. The .indication and .confirm primitives are passed from the MAC to the CL to indicate an internal MAC event that is significant to the CL. This event may be logically related to a remote service request or may be caused by an event internal to the local MAC. The .response primitive is passed from the CL to the MAC to provide a response to a .indication primitive. Thus, the four primitives are used in pairs, the pair .request and .confirm and the pair .indication and .response. This is shown in Figures 8-38 to 8-41.

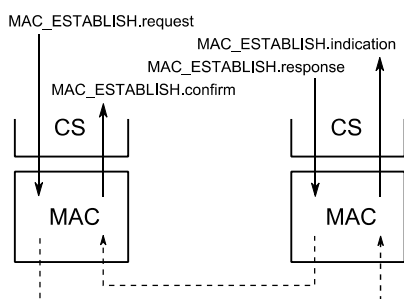


Figure 8-38 – Establishment of a connection




Figure 8-39 – Failed establishment of a connection

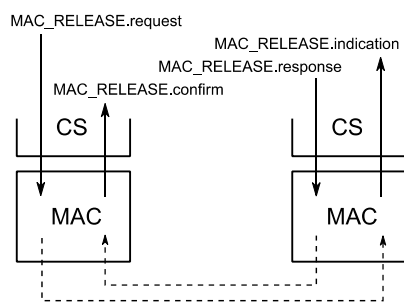


Figure 8-40 – Release of a connection




Figure 8-41 – Transfer of data

Table 8-27 represents the list of available primitives in the MAC-SAP:

Table 8-27 – List of MAC primitives

Service node primitives	Base node primitives
MAC_ESTABLISH.request	MAC_ESTABLISH.request
MAC_ESTABLISH.indication	MAC_ESTABLISH.indication
MAC_ESTABLISH.response	MAC_ESTABLISH.response
MAC_ESTABLISH.confirm	MAC_ESTABLISH.confirm
MAC_RELEASE.request	MAC_RELEASE.request
MAC_RELEASE.indication	MAC_RELEASE.indication
MAC_RELEASE.response	MAC_RELEASE.response
MAC_RELEASE.confirm	MAC_RELEASE.confirm
MAC_JOIN.request	MAC_JOIN.request
MAC_JOIN.Response	MAC_JOIN.response
MAC_JOIN.indication	MAC_JOIN.indication
MAC_JOIN.confirm	MAC_JOIN.confirm
MAC_LEAVE.request	MAC_LEAVE.request

Table 8-27 – List of MAC primitives

Service node primitives	Base node primitives
MAC_LEAVE.indication	MAC_LEAVE.indication
MAC_LEAVE.confirm	MAC_LEAVE.confirm
MAC_DATA.request	MAC_REDIRECT.response
MAC_DATA.confirm	MAC_DATA.request
MAC_DATA.indication	MAC_DATA.confirm
	MAC_DATA.indication

8.5.2 Service node and base node signalling primitives

8.5.2.1 General

The following subclauses describe primitives which are available in both the service node and base node MAC-SAP. These are signalling primitives only and they are used for establishing and releasing MAC connections.

8.5.2.2 MAC_ESTABLISH

8.5.2.2.1 General

The MAC_ESTABLISH primitives are used to manage a connection establishment.

8.5.2.2.2 MAC_ESTABLISH.request

The MAC_ESTABLISH.request primitive is passed to the MAC layer entity to request the connection establishment.

The semantics of this primitive are as follows:

MAC_ESTABLISH.request{EUI-48, Type, Data, DataLength, ARQ, CfBytes}

The EUI-48 parameter of this primitive is used to specify the address of the node to which this connection will be addressed. The MAC will internally transfer this to an address used by the MAC layer. When the CL of a service node wishes to connect to the base node, it uses the EUI-48 00:00:00:00:00:00. However, when the CL of a service node wishes to connect to another service node on the subnetwork, it uses the EUI-48 of that service node. This will then trigger a direct connection establishment. However, whether a normal or a directed connection is established is transparent to the service node MAC SAP. As the EUI-48 of the base node is the SNA, the connection could also be requested from the base node using the SNA.

The *Type* parameter is an identifier used to define the type of the convergence layer that should be used for this connection (see Annex C). This parameter is 1 byte long and will be transmitted in the CON.TYPE field of the connection request.

The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the local CL and the remote CL. This parameter will be transmitted in the CON.DATA field of the connection request.

The *DataLength* parameter is the length of the *Data* parameter in bytes.

The *ARQ* parameter indicates whether or not the ARQ mechanism should be used for this connection. It is a Boolean type with a value of true indicating that ARQ will be used.

The *CfBytes* parameter is used to indicate whether or not the connection should use the contention or contention-free channel access scheme. When *CfBytes* is zero, contention-based access should be used. When *CfBytes* is not zero, it indicates how many bytes per frame should be allocated to the connection using CFP packets.

8.5.2.2.3 MAC_ESTABLISH.indication

The MAC_ESTABLISH.indication is passed from the MAC layer to indicate that a connection establishment was initiated by a remote node.

The semantics of this primitive are as follows:

MAC_ESTABLISH.indication{ConHandle, EUI-48, Type, Data, DataLength, CfBytes}

The *ConHandle* is a unique identifier interchanged to uniquely identify the connection being indicated. It only has a valid meaning in the MAC SAP, and is used to have a reference to this connection between different primitives.

The *EUI-48* parameter indicates which device on the subnetwork wishes to establish a connection.

The *Type* parameter is an identifier used to define the type of the convergence layer that should be used for this connection. This parameter is 1 byte long and it is received in the CON.TYPE field of the connection request.

The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote CL and the local CL. This parameter is received in the CON.DATA field of the connection request.

The *DataLength* parameter is the length of the data parameter in bytes.

The *CfBytes* parameter is used to indicate if the connection should use the contention or contention-free channel access scheme. When *CfBytes* is zero, contention-based access will be used. When *CfBytes* is not zero, it indicates how many bytes per frame the connection would like to be allocated.

8.5.2.2.4 MAC_ESTABLISH.response

The MAC_ESTABLISH.response is passed to the MAC layer to respond with a MAC_ESTABLISH.indication.

The semantics of this primitive are as follows:

MAC_ESTABLISH.response{ConHandle, Answer, Data, DataLength}

The *ConHandle* parameter is the same as the one that was received in the MAC_ESTABLISH.indication.

The *Answer* parameter is used to notify the MAC of the action to be taken for this connection establishment. This parameter may have one of the values in Table 8-28.

The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote CL and the local CL. This parameter is received in the CON.DATA field of the connection response.

The *DataLength* parameter is the length of the data parameter in bytes.

Data may be passed to the caller even when the connection is rejected, i.e., answer has the value 1. The data may then optionally contain more information as to why the connection was rejected.

Table 8-28 – Values of the answer parameter in MAC_ESTABLISH.response primitive

Answer	Description
Accept = 0	The connection establishment is accepted.
Reject = 1	The connection establishment is rejected.

8.5.2.2.5 MAC_ESTABLISH.confirm

The MAC_ESTABLISH.confirm is passed from the MAC layer as the remote answer to a MAC_ESTABLISH.request.

The semantics of this primitive are as follows:

MAC_ESTABLISH.confirm{*ConHandle*, *Result*, *EUI-48*, *Type*, *Data*, *DataLength*}

The *ConHandle* is a unique identifier to uniquely identify the connection being indicated. It has a valid meaning only in the MAC SAP, used to have a reference to this connection between different primitives. The value is only valid if the *Result* parameter is 0.

The *Result* parameter indicates the result of the connection establishment process. It may have one of the values in Table 8-29.

The *EUI-48* parameter indicates which device on the subnetwork wishes to establish a connection.

The *Type* parameter is an identifier used to define the type of the convergence layer that should be used for this connection. This parameter is 1 byte long and it is received in the CON.TYPE field of the connection request.

The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote CL and the local CL. This parameter is received in the CON.DATA field of the connection response.

The *DataLength* parameter is the length of the data parameter in bytes.

Data may be passed to the caller even when the connection is rejected, i.e., the result has the value 1. The data may then optionally contain more information as to why the connection was rejected.

Table 8-29 – Values of the result parameter in MAC_ESTABLISH.confirm primitive

Result	Description
Success = 0	The connection establishment was successful.
Reject = 1	The connection establishment failed because it was rejected by the remote node.
Timeout = 2	The connection establishment process timed out.
No bandwidth = 3	There is insufficient available bandwidth to accept this contention-free connection.
No Such Device = 4	A device with the destination address cannot be found.
Redirect failed = 5	The base node attempted to perform a redirect which failed.
Not Registered = 6	The service node is not registered.
No More LCIDs = 7	All available LCIDs have been allocated.

8.5.2.3 MAC_RELEASE

8.5.2.3.1 General

The MAC_RELEASE primitives are used to release a connection.

8.5.2.3.2 MAC_RELEASE.request

The MAC_RELEASE.request is a primitive used to initiate the release process of a connection.

The semantics of this primitive are as follows:

MAC_RELEASE.request{ConHandle}

The *ConHandle* parameter specifies the connection to be released. This handle is the one that was obtained during the MAC_ESTABLISH primitives.

8.5.2.3.3 MAC_RELEASE.indication

The MAC_RELEASE.indication is a primitive used to indicate that a connection is being released. It may be released because of a remote operation or because of a connectivity problem.

The semantics of this primitive are as follows:

MAC_RELEASE.indication{ConHandle, Reason}

The *ConHandle* parameter specifies the connection being released. This handle is the one that was obtained during the MAC_ESTABLISH primitives.

The *Reason* parameter may have one of the values given in Table 8-30.

Table 8-30 – Values of the reason parameter in MAC_RELEASE.indication primitive

Reason	Description
Success = 0	The connection release was initiated by a remote service.
Error = 1	The connection was released because of a connectivity problem.

8.5.2.3.4 MAC_RELEASE.response

The MAC_RELEASE.response is a primitive used to respond to a connection release process.

The semantics of this primitive are as follows:

MAC_RELEASE.response{ConHandle, Answer}

The *ConHandle* parameter specifies the connection being released. This handle is the one that was obtained during the MAC_ESTABLISH primitives.

The *Answer* parameter may have one of the values given in Table 8-31. This parameter may not have the value "*Reject* = 1" because a connection release process cannot be rejected.

Table 8-31 – Values of the answer parameter in MAC_RELEASE.response primitive

Answer	Description
Accept = 0	The connection release is accepted.

After sending the MAC_RELEASE.response the ConHandle is no longer valid and should not be used.

8.5.2.3.5 MAC_RELEASE.confirm

The MAC_RELEASE.confirm primitive is used to confirm that the connection release process has finished.

The semantics of this primitive are as follows:

MAC_RELEASE.confirm{ConHandle, Result}

The *ConHandle* parameter specifies the connection released. This handle is the one that was obtained during the MAC_ESTABLISH primitives.

The *Result* parameter may have one of the values given in Table 8-32.

Table 8-32 – Values of the result parameter in MAC_RELEASE.confirm primitive

Result	Description
Success = 0	The connection release was successful.
Timeout = 2	The connection release process timed out.
Not Registered = 6	The service node is no longer registered.

After the reception of the MAC_RELEASE.confirm the ConHandle is no longer valid and should not be used.

8.5.2.4 MAC_JOIN

8.5.2.4.1 General

The MAC_JOIN primitives are used to join to a broadcast or multicast connection and allow the reception of such packets.

8.5.2.4.2 MAC_JOIN.request

The MAC_JOIN.request primitive is used:

- by all nodes to join broadcast traffic of a specific CL and start receiving these packets;
- by service nodes to join a particular multicast group;
- by the base node to invite a service node to join a particular multicast group.

Depending on which device makes the join-request, this SAP can have two different variants. The first variant shall be used on base nodes and the second on service nodes.

The semantics of this primitive are as follows:

MAC_JOIN.request{Broadcast, ConHandle, EUI-48, Type, Data, DataLength}

MAC_JOIN.request(Broadcast, Type, Data, Datalength)

The *Broadcast* parameter specifies whether the JOIN operation is being performed for a broadcast connection or for a multicast operation. It should be 1 for a broadcast operation and 0 for a multicast operation. In the case of broadcast operation, EUI-48, Data, DataLength are not used.

ConHandle indicates the handle to be used with for this multicast join. In case of first join request for a new multicast group, ConHandle will be set to 0. For any subsequent EUI additions to an existing multicast group, ConHandle will serve as index to respective multicast group.

The *EUI-48* parameter is used by the base node to specify the address of the node to which this join request will be addressed. The MAC will internally transfer this to an address used by the MAC layer. When the CL of a service node initiates the request, it uses the EUI-48 00:00:00:00:00:00.

The *Type* parameter defines the type of the convergence layer that will send/receive the data packets. This parameter is 1 byte long and will be transmitted in the MUL.TYPE field of the join request.

The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote CL and the local CL. This parameter is received in the MUL.DATA field of the connection request. In case the CL supports several multicast groups, this Data parameter will be used to uniquely identify the group

The *DataLength* parameter is the length of the Data parameter in bytes.

If "Broadcast" is 1, the MAC will immediately issue a MAC_JOIN.confirm primitive, as it does not need to perform any end-to-end operation. For a multicast operation the MAC_JOIN.confirm is only sent once signalling with the uplink service node/base node is complete.

8.5.2.4.3 MAC_JOIN.confirm

The MAC_JOIN.confirm primitive is received to confirm that the MAC_JOIN.request operation has finished.

The semantics of this primitive are as follows:

MAC_JOIN.confirm{*ConHandle*, *Result*}

The *ConHandle* is a unique identifier to uniquely identify the connection being indicated. It has a valid meaning only in the MAC SAP, used to have a reference to this connection between different primitives. The value is only valid if the result parameter is 0. When the MAC receives packets on this connection, they will be passed upwards using the MAC_DATA.indication primitive with this *ConHandle*.

The *Result* parameter indicates the result of multicast group join process. It may have one of the values given in Table 8-33.

Table 8-33 – Values of the Result parameter in MAC_JOIN.confirm primitive

Result	Description
Success = 0	The connection establishment was successful.
Reject = 1	The connection establishment failed because it was rejected by the upstream service node/base node.
Timeout = 2	The connection establishment process timed out.

8.5.2.4.4 MAC_JOIN.indication

On the base node, the MAC_JOIN.indication is passed from the MAC layer to indicate that a multicast group join was initiated by a service node. On a service node, it is used to indicate that the base node is inviting to join a multicast group.

Depending on the device type, this primitive shall have two variants. The first variant below shall be used in base nodes and the second variant is for service nodes:

MAC_JOIN.indication{*ConHandle*, *EUI-48*, *Type*, *Data*, *DataLength*}

MAC_JOIN.indication(*ConHandle*, *Type*, *Data*, *DataLen*)

The *ConHandle* is a unique identifier interchanged to uniquely identify the multicast group being indicated. It has a valid meaning only in the MAC SAP, used to have a reference to this connection between different primitives.

The *EUI-48* parameter indicates which device on the subnetwork wishes to establish a connection.

The *Type* parameter is an identifier used to define the type of the convergence layer that should be used for this request. This parameter is 1 byte long and it is received in the MUL.TYPE field of the connection request.

The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote CL and the local CL. This parameter is received in the MUL.DATA field of the connection request.

The *DataLength* parameter is the length of the Data parameter in bytes.

8.5.2.4.5 MAC_JOIN.response

The MAC_JOIN.response is passed to the MAC layer to respond with a MAC_JOIN.indication. Depending on the device type, this primitive could have either of the two forms given below. The first one shall be used in service node and the second one in base node implementations.

The semantics of this primitive are as follows:

MAC_JOIN.response{ConHandle, Answer}

MAC_JOIN.response (ConHandle, EUI, Answer)

The *ConHandle* parameter is the same as the one that was received in the MAC_JOIN.indication.

EUI is the EUI-48 of the service node that requested the multicast group join.

The *Answer* parameter is used to notify the MAC of the action to be taken for this join request. This parameter may have one of the values depicted below.

Table 8-34 – Values of the answer parameter in MAC_ESTABLISH.response primitive

Answer	Description
Accept = 0	The multicast group join is accepted.
Reject = 1	The multicast group join is rejected.

8.5.2.5 MAC_LEAVE

8.5.2.5.1 General

The MAC_LEAVE primitives are used to leave a broadcast or multicast connection.

8.5.2.5.2 MAC_LEAVE.request

The MAC_LEAVE.request primitive is used to leave multicast or broadcast traffic. Depending on the device type, this primitive could have either of the two forms given below. The first one shall be used in service node and the second one in base node implementations.

The semantics of this primitive are as follows:

MAC_LEAVE.request{ConHandle}

MAC_LEAVE.request{ConHandle, EUI}

The *ConHandle* parameter specifies the connection to be left. This handle is the one that was obtained during the MAC_JOIN primitives.

EUI is the EUI-48 of the service node that has requested to be removed from the multicast group.

8.5.2.5.3 MAC_LEAVE.confirm

The MAC_LEAVE.confirm primitive is received to confirm that the MAC_LEAVE.request operation has finished.

The semantics of this primitive are as follows:

MAC_LEAVE.confirm{ConHandle, Result}

The *ConHandle* parameter specifies the connection released. This handle is the one that was obtained during the MAC_JOIN primitives.

The *Result* parameter may have one of the values in Table 8-35.

Table 8-35 – Values of the result parameter in MAC_LEAVE.confirm primitive

Result	Description
Success = 0	The connection leave was successful.
Timeout = 2	The connection leave process timed out.

After the reception of the MAC_LEAVE.confirm, the ConHandle is no longer valid and should not be used.

8.5.2.5.4 MAC_LEAVE.indication

The MAC_LEAVE.indication primitive is used to leave multicast or broadcast traffic. Depending on the device type, this primitive could have either of the two forms given below. The first one shall be used in service node and the second one in base node implementations.

The semantics of this primitive are as follows:

MAC_LEAVE.indication{ConHandle}

MAC_LEAVE.indication{ConHandle, EUI}

The *ConHandle* parameter is the same as that received in MAC_JOIN.confirm or MAC_JOIN.indication. This handle is the one that was obtained during the MAC_JOIN primitives.

EUI is the EUI-48 of the service node that has requested to be removed from the multicast group.

8.5.3 Base node signalling primitives

8.5.3.1 General

This clause specifies MAC-SAP primitives that are only available in the base node.

8.5.3.2 MAC_REDIRECT.response

The MAC_REDIRECT.response primitive is used to answer to a MAC_ESTABLISH.indication and redirects the connection from the base node to another service node on the subnetwork.

The semantics of this primitive are as follows:

MAC_REDIRECT.reponse{ConHandle, EUI-48, Data, DataLength}

The *ConHandle* is the one passed in the MAC_ESTABLISH.indication primitive to which it is replying.

EUI-48 indicates the service node to which this connection establishment should be forwarded. The base node should perform a direct connection set-up between the source of the connection establishment and the service node indicated by *EUI-48*.

The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote CL and the base node CL. This parameter is received in the CON.DATA field of the connection request.

The *DataLength* parameter is the length of the data parameter in bytes.

Once this primitive has been used, the ConHandle is no longer valid.

8.5.4 Service and base nodes data primitives

8.5.4.1 General

The following subclauses describe how a service node or base node passes data between the convergence layer and the MAC layer.

8.5.4.2 MAC_DATA.request

The MAC_DATA.request primitive is used to initiate the transmission process of data over a connection.

The semantics of the primitive are as follows:

MAC_DATA.request{ConHandle, Data, DataLength, Priority}

The *ConHandle* parameter specifies the connection to be used for the data transmission. This handle is the one that was obtained during the connection establishment primitives.

The *Data* parameter is a buffer of octets that contains the CL data to be transmitted through this connection.

The *DataLength* parameter is the length of the data parameter in octets.

Priority indicates the priority of the data to be sent when using the CSMA access scheme, i.e., the parameter only has meaning when the connection was established with *CfBytes* = 0.

8.5.4.3 MAC_DATA.confirm

The MAC_DATA.confirm primitive is used to confirm that the transmission process of the data has completed.

The semantics of the primitive are as follows:

MAC_DATA.confirm{ConHandle, Data, Result}

The *ConHandle* parameter specifies the connection that was used for data transmission. This handle is the one that was obtained during the connection establishment primitives.

The *Data* parameter is a buffer of octets that contains the CL data to be transmitted through this connection.

The *Result* parameter indicates the result of the transmission. This can take one of the values given in Table 8-36.

Table 8-36 – Values of the result parameter in MAC_DATA.confirm primitive

Result	Description
Success = 0	The send was successful.
Timeout = 2	The send process timed out.

8.5.4.4 MAC_DATA.indication

The MAC_DATA.indication primitive notifies the reception of data through a connection to the CL.

The semantics of the primitive are as follows:

MAC_DATA.indication{ConHandle, Data, DataLength}

The *ConHandle* parameter specifies the connection where the data was received. This handle is the one that was obtained during the connection establishment primitives.

The *Data* parameter is a buffer of octets that contains the CL data received through this connection.

The *DataLength* parameter is the length of the *Data* parameter in octets.

8.5.5 MAC layer management entity SAPs

8.5.5.1 General

The following primitives are all optional.

The aim is to allow an external management entity to control registration and promotion of the service node, and demotion and unregistration of a service node. The MAC layer would normally perform this automatically, however, in some situations/applications it could be advantageous if this could be externally controlled. Indications are also defined so that an external entity can monitor the status of the MAC.

8.5.5.2 MLME_REGISTER

8.5.5.2.1 General

The MLME_REGISTER primitives are used to perform registration and to indicate when registration has been performed.

8.5.5.2.2 MLME_REGISTER.request

The MLME_REGISTER.request primitive is used to trigger the registration process to a subnetwork through a specific switch node. This primitive may be used for enforcing the selection of a specific switch node that may not necessarily be used if the selection is left automatic. The base node MLME function does not export this primitive.

The semantics of the primitive could be either of the following:

$$MLME_REGISTER.request\{\}$$

Invoking this primitive without any parameter simply invokes the registration process in the MAC and leaves the selection of the subnetwork and switch node to MAC algorithms. Using this primitive enables the MAC to perform fully automatic registration if such a mode is implemented in the MAC.

$$MLME_REGISTER.request\{SNA\}$$

The *SNA* parameter specifies the subnetwork to which registration should be performed. Invoking the primitive in this format commands the MAC to only register to the specified subnetwork.

$$MLME_REGISTER.request\{SID\}$$

The *SID* parameter is the SID (switch identifier) of the switch node through which registration needs to be performed. Invoking the primitive in this format commands the MAC to register only to the specified switch node.

8.5.5.2.3 MLME_REGISTER.confirm

The MLME_REGISTER.confirm primitive is used to confirm the status of completion of the registration process that was initiated by an earlier invocation of the corresponding request primitive. The base node MLME function does not export this primitive.

The semantics of the primitive are as follows:

$$MLME_REGISTER.confirm\{Result, SNA, SID\}$$

The *Result* parameter indicates the result of the registration. This can take one of the values given in Table 8-37.

Table 8-37 – Values of the result parameter in MLME_REGISTER.confirm primitive

Result	Description
Done = 0	Registration to a specified SNA through a specified switch is completed successfully.
Timeout =2	Registration request timed out.
Rejected=1	Registration request is rejected by the base node of a specified SNA.
NoSNA=8	Specified SNA is not within range.
NoSwitch=9	Switch node with specified EUI-48 is not within range.

The *SNA* parameter specifies the subnetwork to which registration is performed. This parameter is of significance only if *Result*=0.

The *SID* parameter is the SID (switch identifier) of the switch node through which registration is performed. This parameter is of significance only if *Result*=0.

8.5.5.2.4 MLME_REGISTER.indication

The MLME_REGISTER.indication primitive is used to indicate a status change in the MAC. The service node is now registered to a subnetwork.

The semantics of the primitive are as follows:

MLME_REGISTER.indication{*SNA*, *SID*}

The *SNA* parameter specifies the subnetwork to which registration is performed.

The *SID* parameter is the SID (switch identifier) of the switch node through which registration is performed.

8.5.5.3 MLME_UNREGISTER

8.5.5.3.1 General

The MLME_UNREGISTER primitives are used to perform deregistration and to indicate when deregistration has been performed.

8.5.5.3.2 MLME_UNREGISTER.request

The MLME_UNREGISTER.request primitive is used to trigger the unregistration process. This primitive may be used by management entities if they require the node to unregister for some reason (e.g., register through another switch node or to another subnetwork). The base node MLME function does not export this primitive.

The semantics of the primitive are as follows:

MLME_UNREGISTER.request{}

8.5.5.3.3 MLME_UNREGISTER.confirm

The MLME_UNREGISTER.confirm primitive is used to confirm the status of the completion of the unregister process initiated by an earlier invocation of the corresponding request primitive. The base node MLME function does not export this primitive.

The semantics of the primitive are as follows:

MLME_UNREGISTER.confirm{*Result*}

The *Result* parameter indicates the result of the registration. This can take one of the values given in Table 8-38.

Table 8-38 – Values of the result parameter in MLME_UNREGISTER.confirm primitive

Result	Description
Done = 0	Unregister process completed successfully.
Timeout =2	Unregister process timed out.
Redundant=10	The node is already in disconnected functional state and does not need to unregister.

On generation of MLME_UNREGISTER.confirm, the MAC layer shall not perform any automatic actions that may invoke the registration process again. In such cases, it is up to the management entity to restart the MAC functionality with appropriate MLME_REGISTER primitives.

8.5.5.3.4 MLME_UNREGISTER.indication

The MLME_UNREGISTER.indication primitive is used to indicate a status change in the MAC. The service node is no longer registered to a subnetwork.

The semantics of the primitive are as follows:

MLME_UNREGISTER.indication{}

8.5.5.4 MLME_PROMOTE

8.5.5.4.1 General

The MLME_PROMOTE primitives are used to perform promotion and to indicate when promotion has been performed.

8.5.5.4.2 MLME_PROMOTE.request

The MLME_PROMOTE.request primitive is used to trigger the promotion process in a service node that is in a terminal functional state. This primitive may be used by management entities to enforce promotion in cases where the node's default functionality does not automatically start the process. Implementations may use such triggered promotions to improve the subnetwork topology from time to time.

The semantics of the primitive are as follows:

MLME_PROMOTE.request{}

The value of PRO.PNA in the promotion message sent to the base node is undefined and implementation-specific.

8.5.5.4.3 MLME_PROMOTE.confirm

The MLME_PROMOTE.confirm primitive is used to confirm the status of completion of a promotion process that was initiated by an earlier invocation of the corresponding request primitive.

The semantics of the primitive are as follows:

MLME_PROMOTE.confirm{Result}

The *Result* parameter indicates the result of the registration. This can take one of the values given in Table 8-39.

Table 8-39 – Values of the result parameter in MLME_PROMOTE.confirm primitive

Result	Description
Done = 0	Node is promoted to switch function successfully.
Timeout =1	Promotion process timed out.
Rejected=2	The base node rejected promotion request.
Redundant=10	This device is already functioning as a switch node.

8.5.5.4.4 MLME_PROMOTE.indication

The MLME_PROMOTE.indication primitive is used to indicate a status change in the MAC. The service node is now operating as a switch.

The semantics of the primitive are as follows:

MLME_PROMOTE.indication{}

8.5.5.5 MLME_DEMOTE**8.5.5.5.1 General**

The MLME_DEMOTE primitives are used to perform demotion and to indicate when demotion has been performed.

8.5.5.5.2 MLME_DEMOTE.request

The MLME_DEMOTE.request primitive is used to trigger a demotion process in a service node that is in a switch functional state. This primitive may be used by management entities to enforce demotion in cases where the node's default functionality does not automatically perform the process.

The semantics of the primitive are as follows:

MLME_DEMOTE.request{}

8.5.5.5.3 MLME_DEMOTE.confirm

The MLME_DEMOTE.confirm primitive is used to confirm the status of completion of a demotion process that was initiated by an earlier invocation of the corresponding request primitive.

The semantics of the primitive are as follows:

MLME_DEMOTE.confirm{Result}

The *Result* parameter indicates the result of the demotion. This can take one of the values given in Table 8-40.

Table 8-40 – Values of the result parameter in MLME_DEMOTE.confirm primitive

Result	Description
Done = 0	Node is demoted to terminal function successfully.
Timeout =1	Demotion process timed out.
Redundant=10	This device is already functioning as a terminal node.

When a demotion has been triggered using the MLME_DEMOTE.request, the terminal will remain demoted.

8.5.5.5.4 MLME_DEMOTE.indication

The MLME_DEMOTE.indication primitive is used to indicate a status change in the MAC. The service node is now operating as a terminal.

The semantics of the primitive are as follows:

MLME_DEMOTE.indication{}

8.5.5.6 MLME_RESET

8.5.5.6.1 General

The MLME_RESET primitives are used to reset the MAC into a known "good" status.

8.5.5.6.2 MLME_RESET.request

The MLME_RESET.request primitive results in the flushing of all transmit and receive buffers and the resetting of all state variables. As a result of invoking of this primitive, a service node will transit from its present functional state to the disconnected functional state.

The semantics of the primitive are as follows:

MLME_RESET.request{}

8.5.5.6.3 MLME_RESET.confirm

The MLME_RESET.confirm primitive is used to confirm the status of the completion of a reset process that was initiated by an earlier invocation of the corresponding request primitive. On the successful completion of the reset process, the MAC entity shall restart all functions starting from the search for a subnetwork (clause 8.3.1).

The semantics of the primitive are as follows:

MLME_RESET.confirm{Result}

The *Result* parameter indicates the result of the reset. This can take one of the values given below.

Table 8-41 – Values of the result parameter in MLME_RESET.confirm primitive

Result	Description
Done = 0	MAC reset completed successfully.
Failed =1	MAC reset failed due to internal implementation reasons.

8.5.5.7 MLME_GET

8.5.5.7.1 General

The MLME_GET primitives are used to retrieve individual values from the MAC, such as statistics.

8.5.5.7.2 MLME_GET.request

The MLME_GET.request queries information about a given PIB attribute.

The semantics of the primitive are as follows:

MLME_GET.request{PIBAttribute}

The PIBAttribute parameter identifies specific attributes as listed in the *Id* fields of tables that list PIB attributes (clause 10.2.3).

8.5.5.7.3 MLME_GET.confirm

The MLME_GET.confirm primitive is generated in response to the corresponding MLME_GET.request primitive.

The semantics of this primitive are as follows:

MLME_GET.confirm{status, PIBAttribute, PIBAttributeValue}

The *status* parameter reports the result of requested information and can have one of the values given in Table 8-42.

Table 8-42 – Values of the status parameter in MLME_GET.confirm primitive

Result	Description
Done = 0	Parameter read successfully.
Failed =1	Parameter read failed due to internal implementation reasons.
BadAttr=11	Specified PIBAttribute is not supported.

The *PIBAttribute* parameter identifies specific attributes as listed in *Id* fields of tables that list PIB attributes (clause 10.2.2).

The *PIBAttributeValue* parameter specifies the value associated with a given *PIBAttribute*.

8.5.5.8 MLME_LIST_GET

8.5.5.8.1 General

The MLME_LIST_GET primitives are used to retrieve a list of values from the MAC.

8.5.5.8.2 MLME_LIST_GET.request

The MLME_LIST_GET.request queries a list of values pertaining to a specific class. This special class of PIB attributes are listed in clause 10.2.3.5.

The semantics of the primitive are as follows:

MLME_LIST_GET.request{PIBListAttribute}

The *PIBListAttribute* parameter identifies a specific list that is requested by the management entity. The possible values of *PIBListAttribute* are listed in clause 10.2.3.5.

8.5.5.8.3 MLME_LIST_GET.confirm

The MLME_LIST_GET.confirm primitive is generated in response to the corresponding MLME_LIST_GET.request primitive.

The semantics of this primitive are as follows:

MLME_LIST_GET.confirm{status, PIBListAttribute, PIBListAttributeValue}

The status parameter reports the result of requested information and can have one of the values given in Table 8-43.

Table 8-43 – Values of the status parameter in MLME_LIST_GET.confirm primitive

Result	Description
Done = 0	Parameter read successfully.
Failed =1	Parameter read failed due to internal implementation reasons.
BadAttr=11	Specified PIBListAttribute is not supported.

The *PIBListAttribute* parameter identifies a specific list as listed in the *Id* field of Table 10-7.

The *PIBListAttributeValue* parameter contains the actual listing associated with a given *PIBListAttribute*.

8.5.5.9 MLME_SET

8.5.5.9.1 General

The MLME_SET primitives are used to set configuration values in the MAC.

8.5.5.9.2 MLME_SET.request

The MLME_SET.requests information about a given PIB attribute.

The semantics of the primitive are as follows:

MLME_SET.request{PIBAttribute, PIBAttributeValue}

The *PIBAttribute* parameter identifies a specific attribute as listed in the *Id* fields of tables that list PIB attributes (clause 10.2.3).

The *PIBAttributeValue* parameter specifies the value associated with given *PIBAttribute*.

8.5.5.9.3 MLME_SET.confirm

The MLME_SET.confirm primitive is generated in response to the corresponding MLME_SET.request primitive.

The semantics of this primitive are as follows:

MLME_SET.confirm{Result}

The *Result* parameter reports the result of requested information and can have one of the values given in Table 8-44.

Table 8-44 – Values of the result parameter in MLME_SET.confirm primitive

Result	Description
Done = 0	Given value successfully set for specified attribute.
Failed =1	Failed to set the given value for specified attribute.
BadAttr=11	Specified PIBAttribute is not supported.
OutOfRange=12	Specified PIBAttributeValue is out of acceptable range.
ReadOnly=13	Specified PIBAttributeValue is read only.

8.6 MAC procedures

8.6.1 Registration

The initial service node start-up (clause 8.3.1) is followed by a registration process. A service node in a disconnected functional state shall transmit a REG control packet to the base node in order to get itself included in the subnetwork. Since no LNID or SID is allocated to a service node at this stage, the PKT.LNID field shall be set to all-ones and the PKT.SID field shall contain the SID of the switch node through which it seeks attachment to the subnetwork.

Base nodes may use a registration request as an authentication mechanism. However this Recommendation does not recommend or forbid any specific authentication mechanism and leaves this choice to implementations.

For all successfully accepted registration requests, the base node shall allocate an LNID that is unique within the domain of the switch node through which the attachment is realized. This LNID shall be indicated in the PKT.LNID field of response (REG_RSP). The assigned LNID, in combination with the SID of the switch node through which the service node is registered, would form the NID of the registering node.

Registration is a three-way process. The REG_RSP shall be acknowledged by the receiving service node with a REG_ACK message.

Figure 8-42 represents a successful registration process and Figure 8-43 shows a registration request that is rejected by the base node. Details on specific fields that distinguish one registration message from the other are given in Table 8-9.

The REG control packet, in all its usage variants, is transmitted unencrypted, but specified fields (REG.SNK and REG.AUK) are encrypted with context-specific encryption keys as explained in clause 8.4.5.3. The encryption of REG.AUK in REG_RSP, its decryption at the receiving end and subsequent encrypted retransmission using a different encryption key authenticates that the REG_ACK is from the intended destination.

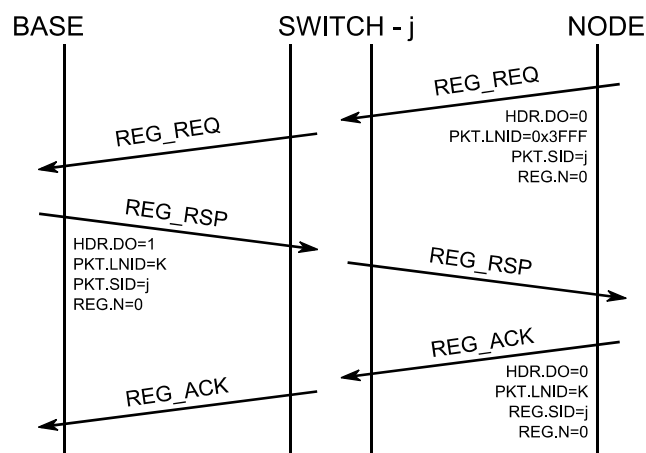


Figure 8-42 – Registration process accepted




Figure 8-43 – Registration process rejected

When assigning an LNID, the base node shall not reuse an LNID released by an unregister process until after $(macMaxCtlReTx + 1) \times macCtlReTxTimer$ seconds, to ensure that all retransmit packets have left the subnetwork. Similarly, the base node shall not reuse an LNID freed by the keep-alive process until T_{keep_alive} seconds have passed, using the last known acknowledged T_{keep_alive} value, or if larger, the last unacknowledged T_{keep_alive} , for the service node using the LNID.

During network start-up where the whole network is powered on at once, there will be considerable contention for the medium. It is recommended, but optional, that randomness is added to the first transmission of REQ_REQ and all subsequent retransmissions. A random delay of maximum duration of 10% of $macCtlReTxTimer$ may be imposed before the first REG_REQ message and a similar random delay of up to 10% of $macCtlReTxTimer$ may be added to each retransmission.

8.6.2 Unregistration process

At any point in time, either the base node or the service node may decide to close an existing registration. This version of the Recommendation does not provide provision for rejecting an unregister request. The service node or base node that receives an unregister request shall acknowledge its receipt and take appropriate actions.

Following a successful unregister, a service node shall move back from its present functional state to a disconnected functional state and the base node may re-use any resources that were reserved for the unregistering node.

Figure 8-44 shows a successful unregister process initiated from a service node and Figure 8-45 shows an unregister process initiated from the base node. Details on specific fields that identify unregister requests in REG control packets are given in Table 8-9




Figure 8-44 – Unregistering process initiated by a terminal node




Figure 8-45 – Unregistering process initiated by the base node

8.6.3 Promotion process

A node that cannot reach any existing switch may send promotion-needed frames so that a terminal can be promoted and begin to switch. During this process, a node that cannot reach any existing switch may send PNPDU's so that a nearby terminal can be promoted and begin to act as a switch. During this process, a terminal will receive PNPDU's and at its discretion, generate PRO_REQ control packets to the base node.

The base node examines the promotion requests during a period of time. It may use the address of the new terminal, provided in the promotion-request packet, to decide whether or not to accept the promotion. It will decide which node shall be promoted, if any, sending a promotion response. The other nodes will not receive any answer to the promotion request to avoid subnetwork saturation. Eventually, the base node may send a rejection if any special situation occurs. If the subnetwork is specially preconfigured, the base node may send terminal node promotion requests directly to a terminal node.

When a terminal node requests promotion, the PRO.NSID field in the PRO_REQ_S message shall be set to all-ones. The PRO.NSID field shall contain an LSID allocated to the promoted node in the PRO_REQ_B message. The acknowledging switch node shall set the PRO.NSID field in its PRO_ACK to the newly allocated LSID. This final PRO_ACK shall be used by intermediate switch nodes to update their switching tables as described in clause 8.3.5.2.

When reusing LSIDs that have been released by a demotion process, the base node should not allocate the LSID until after $(macMaxCtlReTx + 1) \times macCtlReTxTimer$ seconds to ensure all retransmit packets that might use that LSID have left the subnetwork. Similarly, the base node shall not reuse an LNID freed by the keep-alive process until T_{keep_alive} seconds have passed, using the last known acknowledged T_{keep_alive} value, or if larger, the last unacknowledged T_{keep_alive} , for the service node using the LNID.

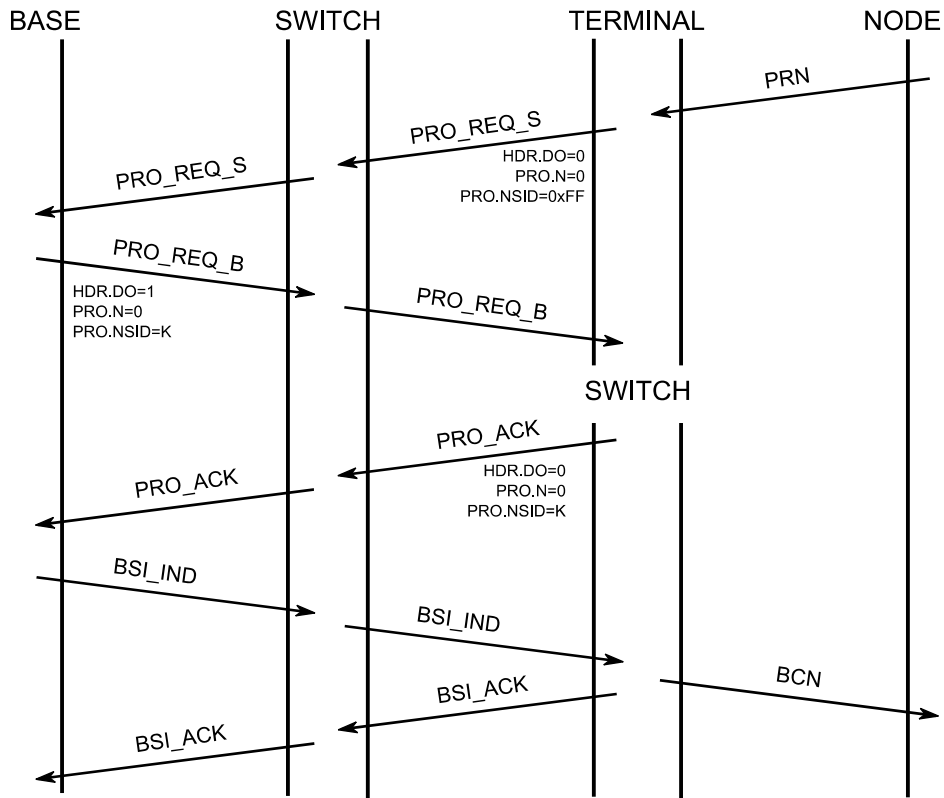


Figure 8-46 – Promotion process initiated by a service node

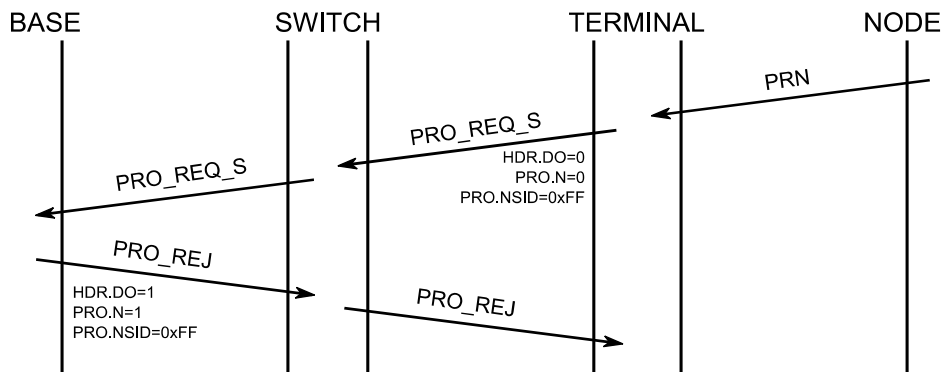


Figure 8-47 – Promotion process rejected by the base node

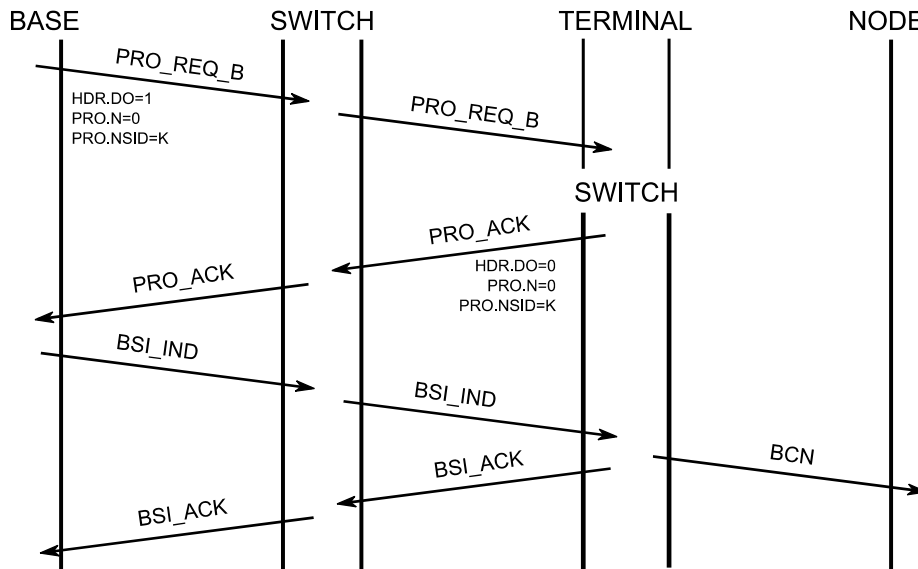


Figure 8-48 – Promotion process initiated by the base node

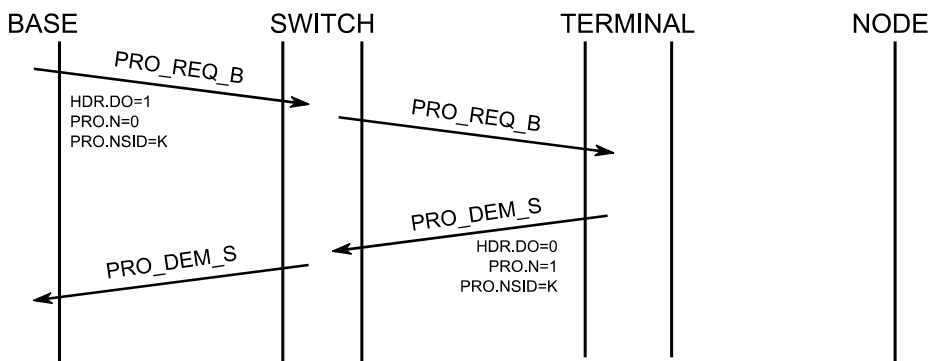


Figure 8-49 – Promotion process rejected by a service node

8.6.4 Demotion process

The base node or a switch node may decide to discontinue a switching function at any time. The demotion process provides such a mechanism. The PRO control packet is used for all demotion transactions.

The PRO.NSID field shall contain the SID of the switch node that is being demoted as part of the demotion transaction. The PRO.PNA field is not used in any demotion process transaction and its contents are not interpreted at either end.

Following the successful completion of a demotion process, a switch node shall immediately stop the transmission of beacons and change from a switch functional state to a terminal functional state. The base node may reallocate the LSID and beacon slot used by the demoted switch after $(macMaxCtlReTx + 1) \times macCtlReTxTimer$ seconds to other terminal nodes requesting promotion.

The present version of this Recommendation does not specify any explicit message to reject a demotion requested by a peer at the other end.

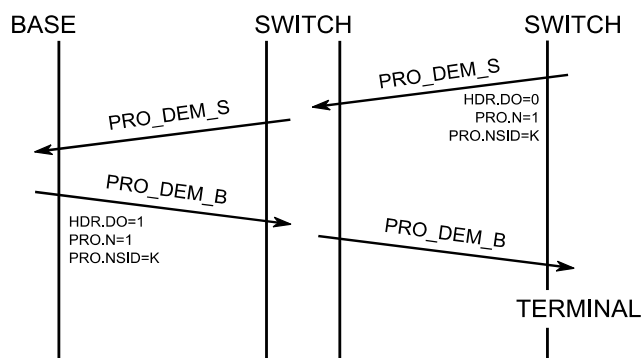


Figure 8-50 – Demotion process initiated by a service node

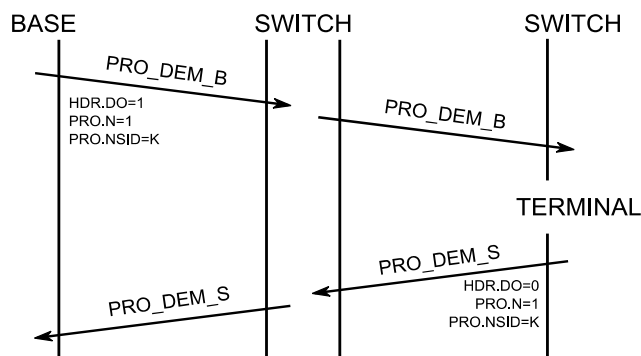


Figure 8-51 – Demotion process initiated by the base node

8.6.5 Keep-alive process

The keep-alive process is used to detect when a service node has left the subnetwork because of changes to the network configuration or because of fatal errors it cannot recover from.

When the service node receives the REG_RSP packet it uses the REG.TIME field to start a timer $T_{\text{keep_alive}}$. For every ALV_B it receives, it restarts this timer using the value from ALV.TIME. It should also send an ALV_S to the base node. If the timer ever expires, the service node assumes it has been unregistered by the base node. The message PRO_REQ does also reset the keep-alive timer to the PRO.TIME value.

Each switch along the path of an ALV_B message should keep a copy of the PRO.TIME and then ALV.TIME for each switch node below it in the tree. When the switch does not receive an ALV_S message from a service node below it for $T_{\text{keep_alive}}$ as defined in PRO.TIME and ALV.TIME it should remove the switch node entry from its switch table. See clause 8.3.5.2 for more information on the switching table. Additionally a switch node may use the REG.TIME and ALV.TIME to also consider every service node registration status and take it into account for the switching table.

For every ALV_S or ALV_B message sent by the base node or service node, the counter ALV.TXCNT should be incremented before the message is sent. This counter is expected to wrap around. For every ALV_B or ALV_S message received by the service node or the base node the counter ALV.RXCNT should be incremented. This counter is also expected to wrap around. These two counters are placed into the ALV_S and ALV_B messages. The base node should keep an ALV.TXCNT and ALV.RXCNT separated counter for each service node. These counters are reset to zero in the registration process.

The algorithm used by the base node to determine when to send ALV_B messages to registered service nodes and how to determine the value ALV.TIME and REG.TIME is not specified here.

8.6.6 Connection management

8.6.6.1 Connection establishment

Connection establishment works end-to-end, connecting the application layers of communicating peers. Owing to the tree topology, most connections in a subnetwork will involve the base node at one end and a service node at the other. However, there may be cases when two service nodes within a subnetwork need to establish connections. Such connections are called direct connections and are described in clause 8.3.6.

All connection establishment messages use the CON control packet. The various control packets types and specific fields that unambiguously identify them are given in Table 8-11.

Each successful connection established on the subnetwork is allocated an LCID. The base node shall allocate an LCID that is unique for a given LNID.

NOTE – Either of the negotiating ends may decide to reject a connection establishment request. The receipt of a connection rejection does not amount to any restrictions on making future connection requests; it may however be advisable.




Figure 8-52 – Connection establishment initiated by a service node




Figure 8-53 – Connection establishment rejected by the base node




Figure 8-54 – Connection establishment initiated by the base node




Figure 8-55 – Connection establishment rejected by a service node

8.6.6.2 Connection closing

Either peer at both ends of a connection may decide to close the connection at any time. The CON control packet is used for all messages exchanged in the process of closing a connection. Only the CON.N field in the CON control packet is relevant in closing an active connection.

A connection closure request from one end is acknowledged by the other end before the connection is considered closed. The present version of this Recommendation does not have any explicit message for rejecting a connection termination requested by a peer at the other end.

Figures 8-56 and 8-57 show message exchange sequences in a connection closing process.

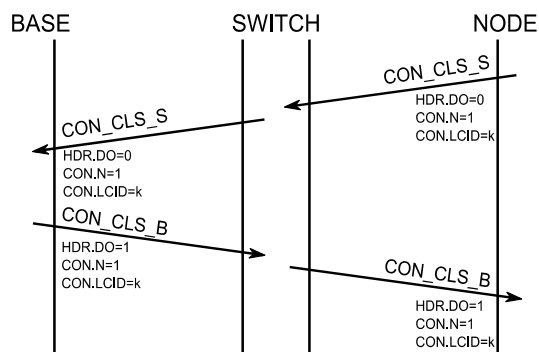


Figure 8-56 – Disconnection initiated by a service node

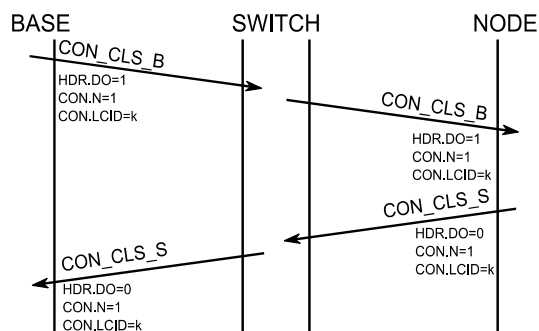


Figure 8-57 – Disconnection initiated by the base node

8.6.7 Multicast group management

8.6.7.1 General

The joining and leaving of a multicast group can be initiated by the base node or the service node. The MUL control packet is used for all messages associated with multicast and the usual retransmit mechanism for control packets is used. These control messages are unicast between the base node and the service node.

8.6.7.2 Group join

Multicast group join may be initiated from either the base node or service node. A device shall not start a new join procedure before an existing join procedure started by itself is completed.

Certain applications may require the base node to selectively invite certain service nodes to join a specific multicast group. In such cases, the base node starts a new group and invites service nodes as required by application.

Successful and failed group joins initiated from the base node are shown in Figures 8-58 and 8-59.

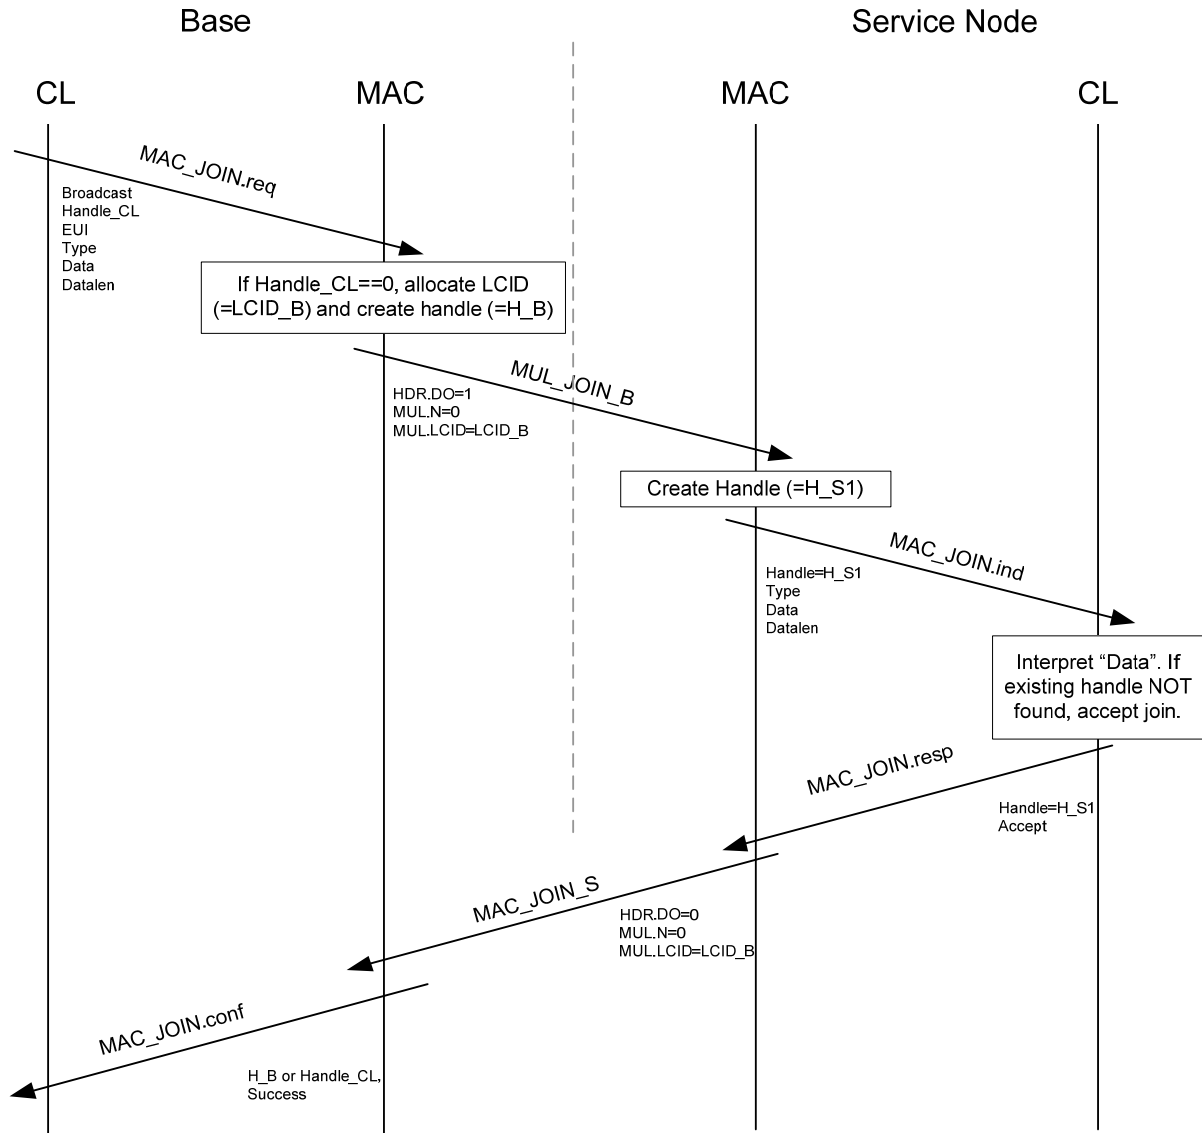


Figure 8-58 – Successful group join initiated by the base node

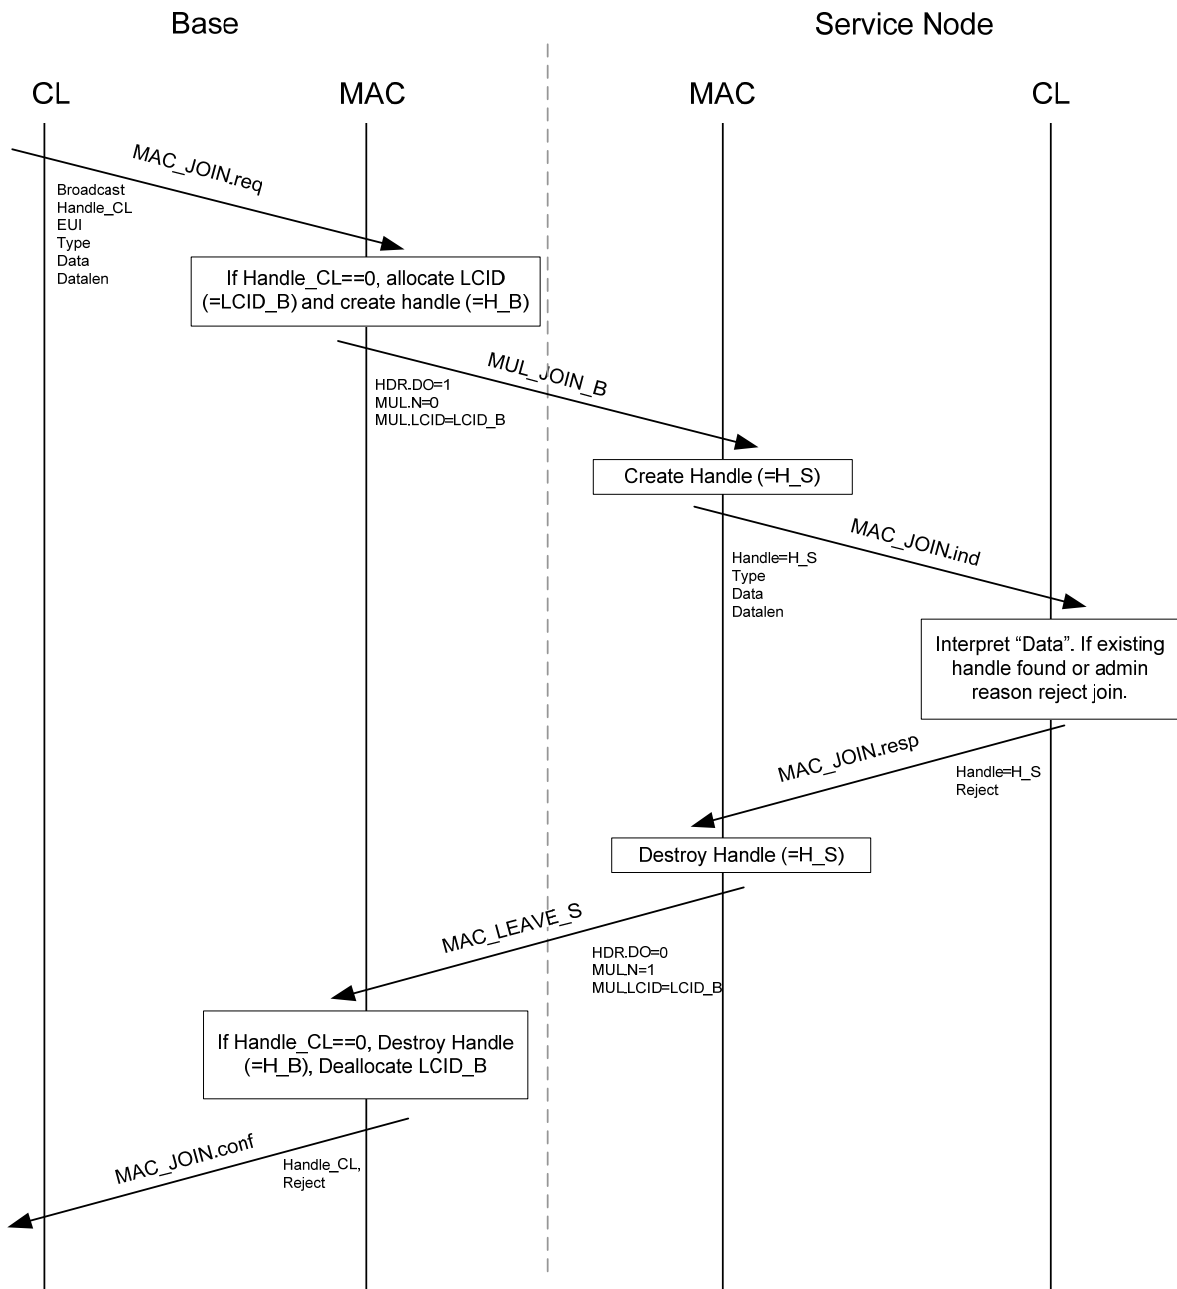


Figure 8-59 – Failed group join initiated by the base node

Successful and failed group joins initiated from the service node are shown in Figures 8-60 and 8-61.

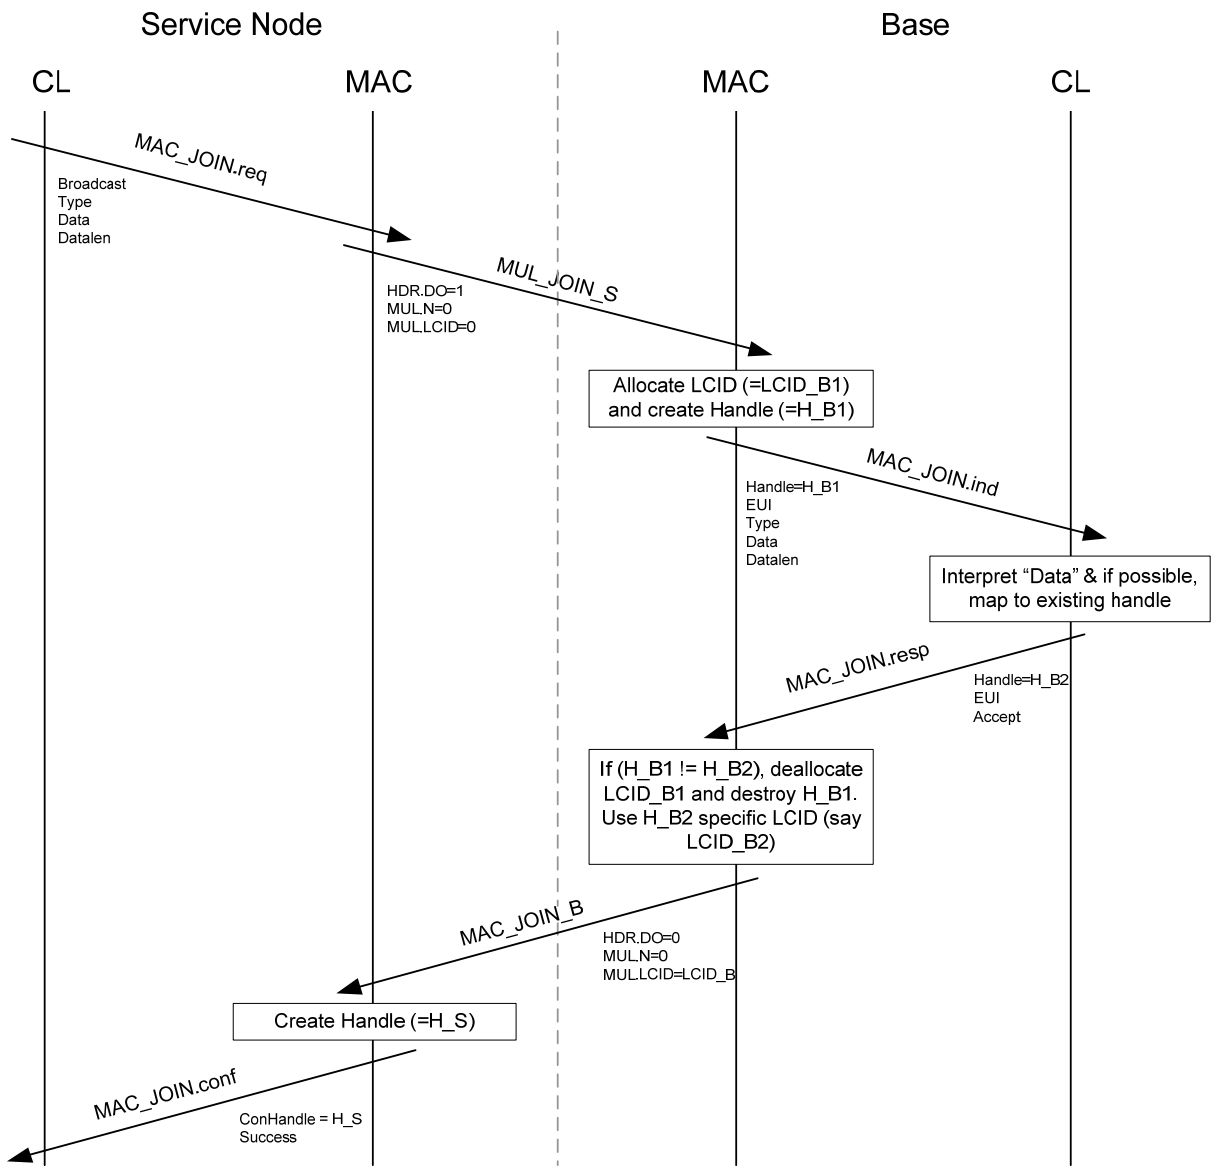


Figure 8-60 – Successful group join initiated by the service node

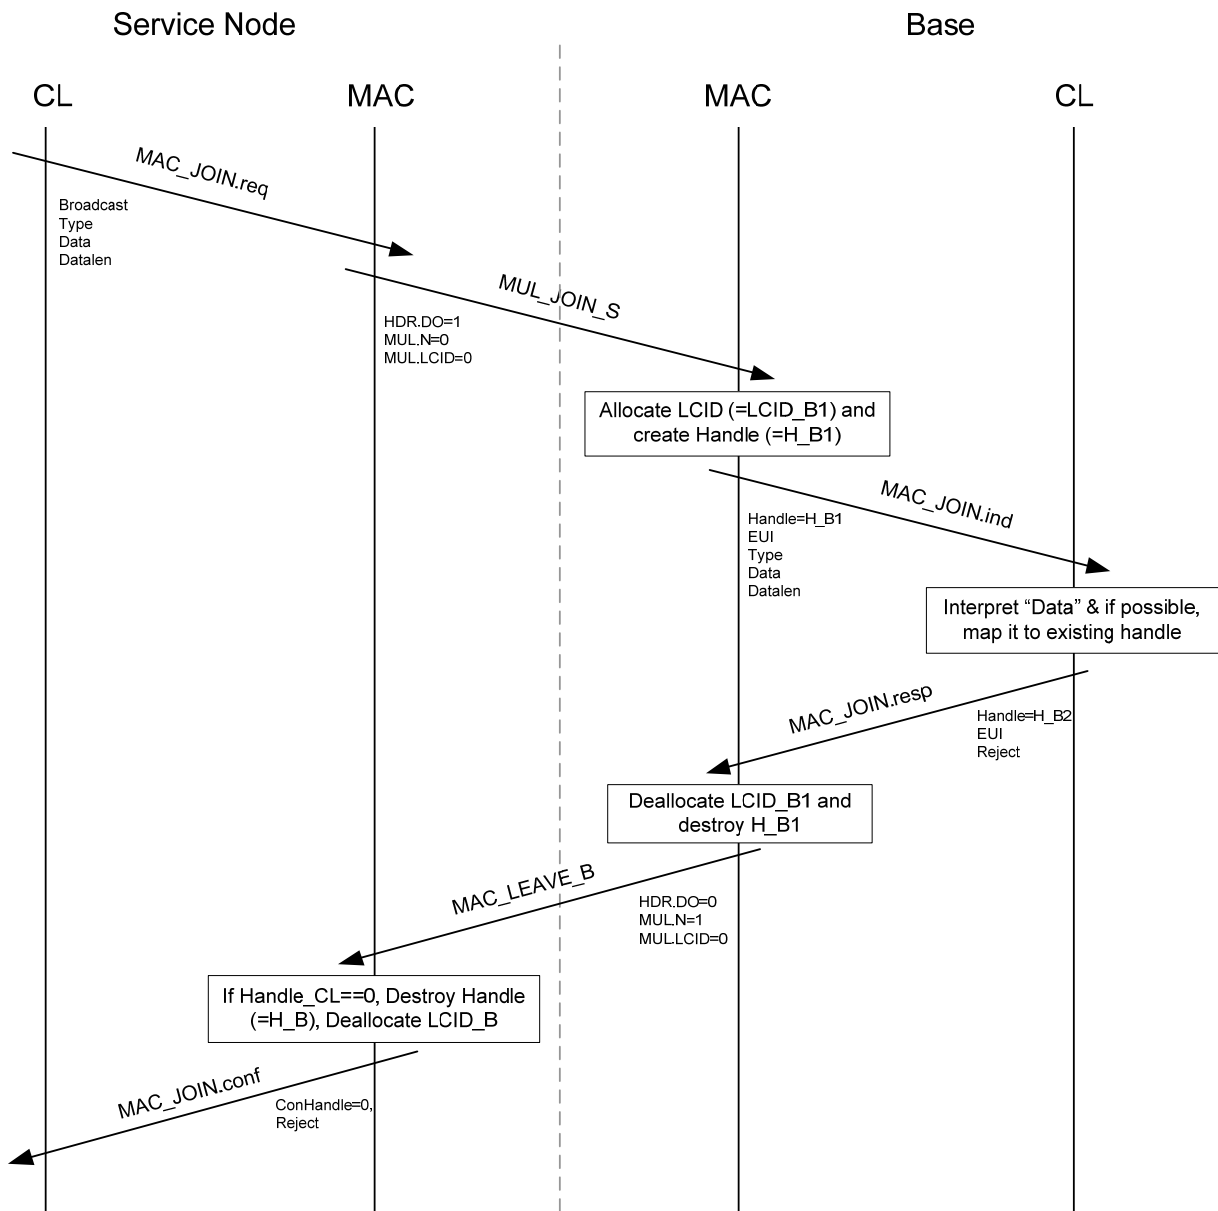


Figure 8-61 – Failed group join initiated by the service node

8.6.7.3 Group leave

Leaving a multicast group operates in the same way as connection removal. Either the base node or service node may decide to leave the group. A notable difference in the group leave process as compared to a group join is that there is no message sequence for rejecting a group leave request.

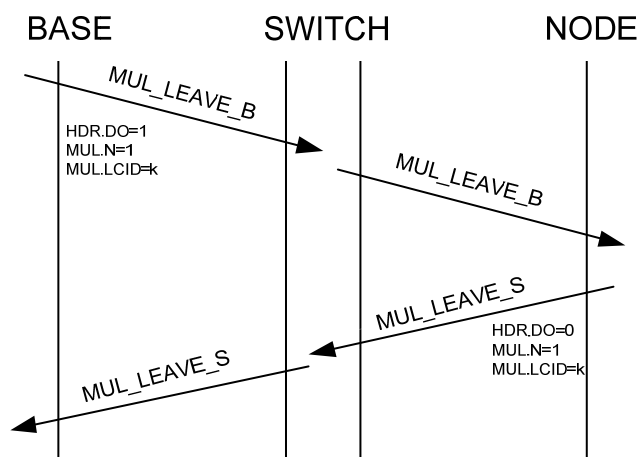


Figure 8-62 – Leave initiated by the service node

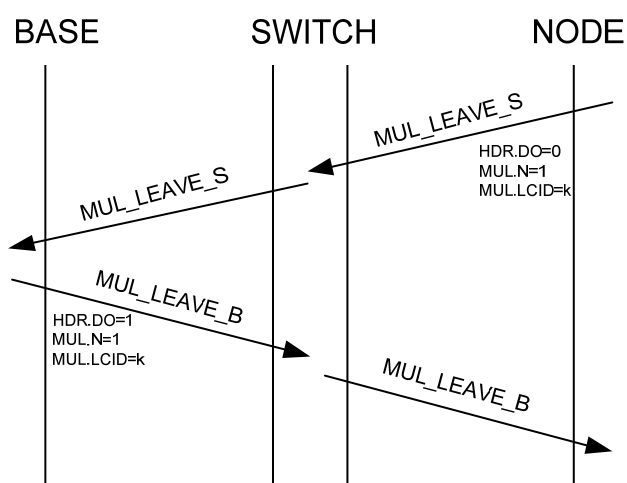


Figure 8-63 – Leave initiated by the base node

8.6.8 PHY robustness management

8.6.8.1 General

The PHY layer has several parameters that affect the performance of the transmission: power transmission, modulation schema (constellation mapping and convolutional encoding). The transmitter needs feedback about the reception quality to adjust its transmission parameters. This feedback is sent using PRM control packets.

8.6.8.2 Use of the PHY robustness notification for detecting robustness needs

There are several sources of information that may be used to detect whether or not the robustness of the PHY is the right one:

- received packets with invalid CRC;
- ARQ retransmissions;
- control packet retransmissions;
- PRM requests sent by other nodes to the same switch node (in the case of node-to-switch notifications);
- PRM responses.

This information may be used to decide when to notify that the robustness of the PHY should be changed. This notification may be performed from a service node to a switch node and from a switch node to a service node. For this purpose, the base node works as the root switch in exactly the same way the other switch nodes do.

8.6.8.3 PHY robustness notification




Figure 8-64 – PHY robustness management requested by the switch node

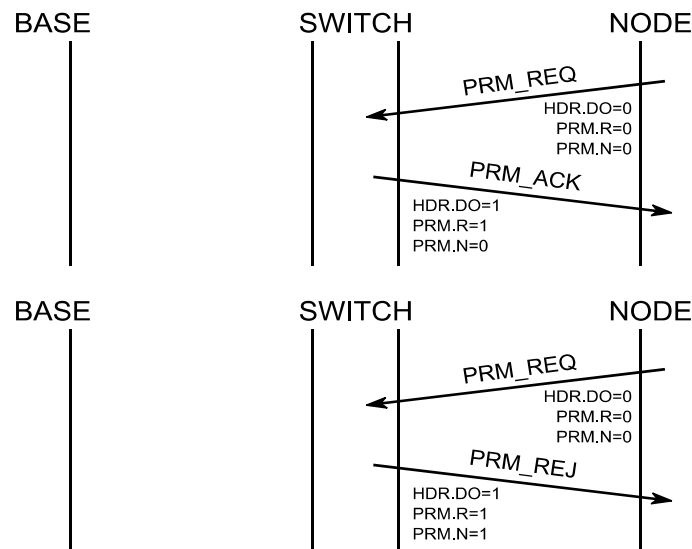


Figure 8-65 – PHY robustness management requested by the service node

8.6.8.4 PHY robustness changing

From the PHY point of view there are several parameters that may be adjusted and which affect the transmission robustness: the transmission power and modulation parameters (convolutional encoding and constellation). As a general rule the following rules should be followed:

- Increase robustness: increase the power and if it is not possible, improve the modulation scheme robustness (reducing throughput).
- Reduce robustness: reduce the modulation scheme robustness (increasing throughput) and if it is not possible, reduce the transmission power.

8.6.9 Channel allocation and de-allocation

Figure 8-66 below shows a successful channel allocation sequence. All channel allocation requests are forwarded to the base node. Note that in order to assure a contention-free channel allocation along the entire path, the base node allocates non-overlapping times to intermediate switch nodes. In a multi-level subnetwork, the base node may also reuse the allocated time at different levels. While reusing the said time, the base node needs to ensure that the levels that use the same time slots have sufficient separation so that there is no possible interference.

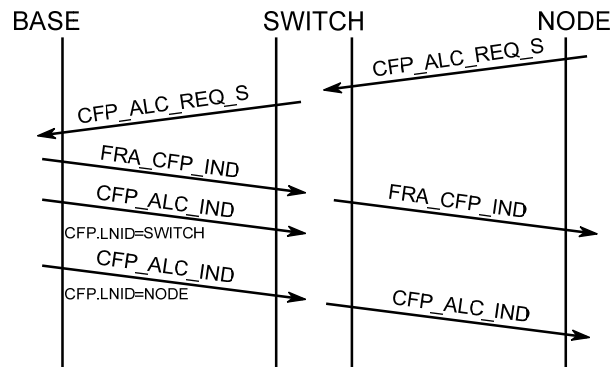


Figure 8-66 – Successful allocation of a CFP period

Figure 8-67 below shows a channel de-allocation request from a terminal device and the resulting confirmation from the base node.

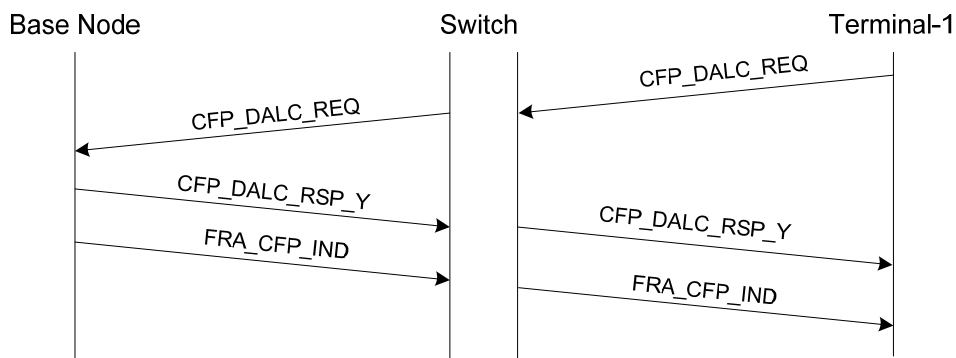


Figure 8-67 – Successful channel de-allocation sequence

Figure 8-68 below shows a sequence of events that may lead to a base node re-allocation contention-free slot to a terminal device that already has slots allocated to it. In this example, a de-allocation request from Terminal-2 resulted in two changes: firstly, in the global frame structure, this change is conveyed to the subnetwork in the FRA_CFP_IND packet; secondly, it is specific to the time slot allocated to Terminal-1 within the CFP.

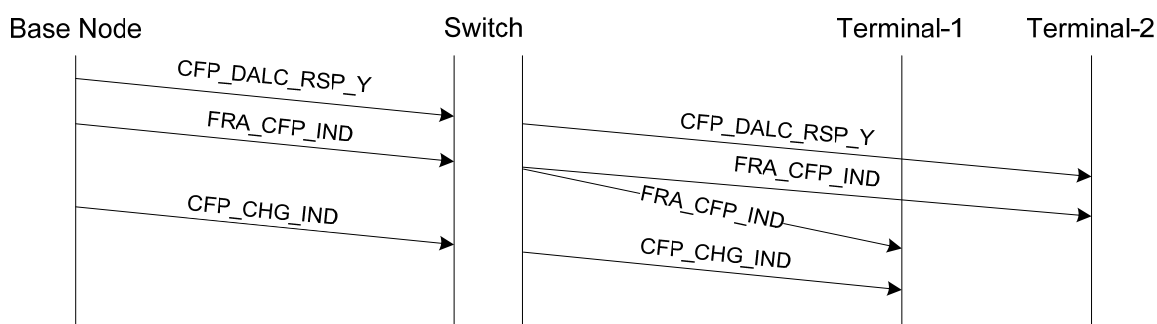


Figure 8-68 – De-allocation of channel to one device results in the change of the CFP allocated to another

8.7 Automatic repeat request (ARQ)

8.7.1 General

Devices complying with this Recommendation may either implement an ARQ scheme as described in this clause or no ARQ at all. This Recommendation provides low-complexity switch and terminal devices that choose not to implement any ARQ mechanism at all.

8.7.2 Initial negotiation

ARQ is a connection property. During the initial connection negotiation, the originating device indicates its preference for ARQ or non-ARQ in the CON.ARQ field. The responding device at the other end can indicate its acceptance or rejection of the ARQ in its response. If both devices agree to use ARQ for the connection, all traffic in the connection will use ARQ for acknowledgements, as described in clause 8.7.3. If the responding device rejects the ARQ in its response, the data flowing through this connection will not use ARQ.

8.7.3 ARQ mechanism

8.7.3.1 General

The ARQ mechanism works between directly connected peers (original source and final destination), as long as both of them support ARQ implementation. This implies that even for a connection between the base node and a terminal (connected via one or more intermediate switch devices), ARQ works on an end-to-end basis. The behaviour of switch nodes in an ARQ-enabled connection is described in clause 8.7.4. When using ARQ, a unique packet identifier is associated with each packet, to aid in acknowledgement. The packet identifier is 6 bits long and can therefore denote 64 distinct packets. ARQ windowing is supported, with a maximum window size of 32 (5 bits), as described in clause 8.7.3.3.

8.7.3.2 ARQ PDU

8.7.3.2.1 General

The ARQ subheader is placed inside the data packets, after the packet header and before the ORIGINAL packet payload.




Figure 8-69 – MAC data PDU with ARQ subheader

For an ARQ PDU, the PKT.LEN field in the packet header will be set as the ARQ subheader length plus the original packet payload length. By doing this, the intermediate switching node can correctly parse the whole PDU length without the knowledge that this PDU is ARQ-enabled, so that it can transparently relay the ARQ PDU based on the addressing information alone.

The ARQ subheader contains a set of bytes; each byte containing different subfields. The most significant bit of each byte, the M bit, indicates if there are more bytes in the ARQ subheader.




Figure 8-70 – ARQ subheader only with the packet id

Figure 8-70 shows an ARQ subheader with the first M bit of 0 and so the subheader is a single byte and contains only the packet ID for the transmitted packet.




Figure 8-71 – ARQ subheader with ARQ.INFO

Figure 8-71 has the M bit in the first byte of the ARQ subheader set, and so the subheader contains multiple bytes. The first byte contains the packet ID of the transmitted packet and then follows the ARQ.INFO which is a list of one or more bytes, where each byte could have one of the following meanings:




Figure 8-72 – ARQ.ACK byte fields




Figure 8-73 – ARQ.WIN byte fields




Figure 8-74 – ARQ.NACK byte fields

If there are multiple packets lost, an ARQ.NACK is sent for each of them, from the first packet lost to the last packet lost. When there are several ARQ.NACK they implicitly acknowledge the packets before the first ARQ.NACK, and the packets in between the ARQ.NACKs. If an ARQ.ACK is present, it must be placed at the end of the ARQ subheader, and should refer to an ARQ.ACKID that is later than any other ARQ.NACKID, if present. If there is at least an ARQ.NACK and an ARQ.ACK they also implicitly acknowledge any packet in the middle between the last ARQ.NACKID and the ARQ.ACK.

For interoperability, a device should be able to receive any well-formed ARQ subheader and should process at least the first ARQ.ACK or ARQ.NACK field.

The subfields have the following meanings as described in Table 8-45.

Table 8-45 – ARQ fields

Field	Description
ARQ.FLUSH	ARQ.FLUSH = 1 If an ACK must be sent immediately. ARQ.FLUSH = 0 If an ACK is not needed.
ARQ.PKTID	The ID of the current packet, if the packet is empty (with no data) this is the ID of the packet that will be sent next.
ARQ.ACKID	The identifier with the next packet expected to be received.
ARQ.WINSIZE	The window size available from the last acknowledged packet. After a connection is established its window is 1.
ARQ.NACKID	Ids of the packets that need to be retransmitted.

8.7.3.2.2 ARQ subheader example

MSB							
ARQ. M = 1	ARQ. FLUSH = 1		ARQ.PKTID = 23	ARQ. M = 1	0	Res	ARQ.WINSIZE = 16
ARQ. M = 1	1		ARQ.NACKID = 45	ARQ. M = 1	1		ARQ.NACKID = 47
ARQ. M = 1	1		ARQ.NACKID = 48	ARQ. M = 1	1		ARQ.NACKID = 52
ARQ. M = 1	1		ARQ.NACKID = 55	ARQ. M = 1	1		ARQ.NACKID = 56
ARQ. M = 1	1		ARQ.NACKID = 57	ARQ. M = 0	0		ARQ.ACKID = 60
LSB							

Figure 8-75 – Example of an ARQ subheader with all the fields present

In this example all the ARQ subheader fields are present. To make it understandable, since both nodes are both transmitters and receivers, the side receiving this header will be called A and the other side transmitting B. The message has the packet ID of 23 if it contains data; otherwise the next data packet to be sent has the packet ID of 23. Since the flush bit is set it needs to be ACKed/NACKed.

B requests the retransmission of packets 45, 47, 48, 52, 55, 56 and 57. ACK = 60, so it has received packets <45, 46, 49, 50, 51, 53, 54, 58 and 59.

The window is 16 and it has received and processed up to packet 44 (first NACK = 45), so A can send all packets <= 60; that is, as well as sending the requested retransmits, it can also send packet ID = 60.

8.7.3.3 Windowing

A new connection between two peer devices starts with an implicit initial receiver window size of 1 and a packet identifier 0. This window size is a limiting case and the transaction (to start with) would behave like a "Stop and Wait" ARQ mechanism.

Upon receipt of an ARQ.WIN, the sender would adapt its window size to *ARQ.WINSIZE*. This buffer size is counted from the first packet completely ACK-ed, so if there is a NACK list and then an ACK the window size defines the number of packets from the first NACK-ed packet that could be sent. If there is just an ACK in the packet (without any NACK) the window size determines the number of packets that can be sent from that ACK.

An *ARQ.WINSIZE* value of 0 may be transmitted back by the receiver to indicate congestion at its end. In such cases, the transmitting end should wait for at least *ARQCongClrTime* before re-transmitting its data.

8.7.3.4 Flow control

The transmitter must manage the ACK sending algorithm by the flush bit; it is up to it having a proper ARQ communication. The receiver is only forced to send ACKs when the transmitter has sent a packet with the flush bit set, although the receiver could send more ACKs even if not forced to do it, because the flow control is only a responsibility of the transmitter.

These are the requisites to be interoperable, but the algorithm is up to the manufacturer. It is strongly recommended to piggyback data-ACK information in outgoing packets, to avoid the transmission of unnecessary packets just for ACK-ing.

8.7.3.5 Algorithm recommendation

No normative algorithm is specified.

8.7.3.6 Usage of ARQ in resource limited devices

Resource-limited devices may have a low memory and simple implementation of ARQ. They may want to use a window of 1 packet. They will work as a "Stop and Wait" mechanism.

The ARQ subheader to be generated may be one of the following:

If there is nothing to acknowledge:




Figure 8-76 – Stop and wait ARQ subheader with only packet ID

If there is something to acknowledge carrying data:



Figure 8-77 – Stop and wait ARQ subheader with an ACK

If there is something to acknowledge but without any data in the packet:



Figure 8-78 – Stop and wait ARQ subheader without data and with an ACK

The ARQ.WINSIZE is not generally transmitted because the window size is already 1 by default, it may only be transmitted to handle congestion and to resume the transmission again.

8.7.4 ARQ packets switching

All switch nodes shall support transparent bridging of ARQ traffic, whether or not they support ARQ for their own transmission and reception. In this mode, switch nodes are not required to buffer the packets of the ARQ connections for retransmission.

Some switch nodes may buffer the packets of the ARQ connections, and perform retransmission in response to NACKs for these packets. The following general principles shall be followed:

- The acknowledged packet identifiers shall have end-to-end coherency.
- The buffering of packets in switch nodes and their retransmissions shall be transparent to the source and destination nodes, i.e., a source or destination node shall not be required to know whether or not an intermediate switch has buffered packets for switched data.

9 Convergence layer

9.1 Overview

Figure 9-1 shows the overall structure of the convergence layer.

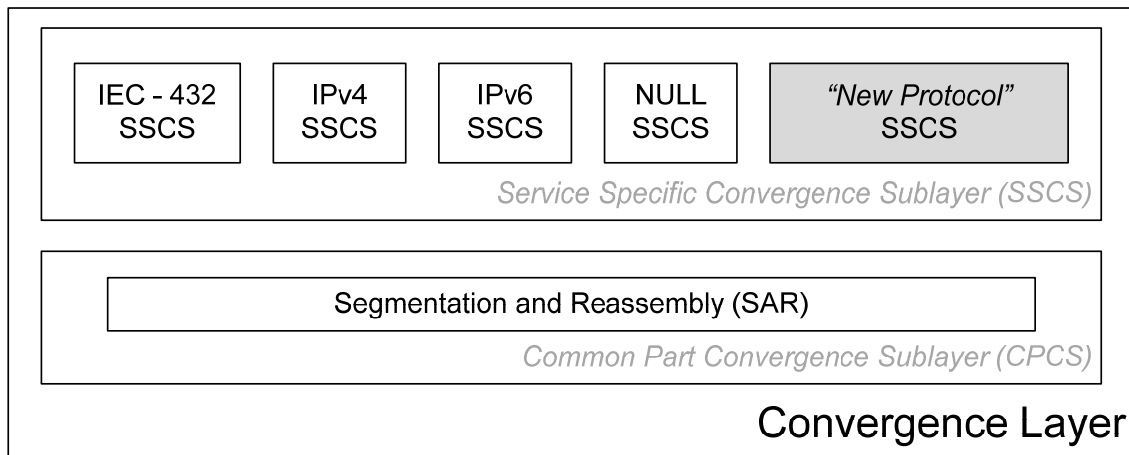


Figure 9-1 – Structure of the convergence layer

The convergence layer is separated into two sublayers. The common part convergence sublayer (CPCS) provides a set of generic services. The service-specific convergence sublayer (SSCS) contains services that are specific to one communication profile. There are several SSCSs, typically one per communication profile, but only one CPCS. The use of CPCS services is optional in that a certain SSCS will use the services it needs from the CPCS, and omit services which are not needed.

9.2 Common part convergence sublayer (CPCS)

9.2.1 General

This Recommendation defines only one CPCS service: segmentation and reassembly (SAR).

9.2.2 Segmentation and reassembly (SAR)

9.2.2.1 General

CPCS SDUs which are larger than 'CIMTUSize-1' bytes are segmented at the CPCS. CPCS SDUs which are equal or smaller than 'CIMTUSize-1' bytes may also optionally be segmented. Segmentation means breaking up a CPCS SDU into smaller parts to be transferred by the MAC layer. At the peer CPCS, the smaller parts (segments) are put back together (i.e., reassembled) to form the complete CPCS SDU. All segments except the last segment of a segmented SDU must be the same size and at most *CIMTUSize* bytes in length. Segments may be decided to be smaller than 'CIMTUSize-1' bytes, e.g., when the channel is poor. The last segment may of course be smaller than 'CIMTUSize-1' bytes.

In order to keep SAR functionality simple, the *CIMTUSize* is a constant value for all possible modulation/coding combinations at the PHY layer. The value of *CIMTUSize* is such that with any modulation/coding combination it is always possible to transmit a single segment in one PPDU. Therefore, there is no need for discovering a specific MTU between peer CPCSs or modifying the SAR configuration for every change in the modulation/coding combination. In order to increase efficiency, a service node which supports packet aggregation may combine multiple segments into one PPDU when communicating with its peer.

Segmentation always adds a 1-byte header to each segment. The first 2 bits of an SAR header identify the type of segment. The semantics of the rest of the header information then depend on the type of segment. The structure of different header types is shown in Figure 9-2 and individual fields are explained in Table 9-1. Not all fields are present in each SAR header. Either SAR.NSEGS or SAR.SEQ is present, but not both.




Figure 9-2 – Segmentation and reassembly headers

Table 9-1 – SAR header fields

Name	Length	Description
SAR.TYPE	2 bits	Type of segment. <ul style="list-style-type: none"> • 0b00: first segment • 0b01: intermediate segment • 0b10: last segment • 0b11: reserved by ITU-T.
SAR.NSEGS	6 bits	'Number of Segments' – 1
SAR.SEQ	6 bits	Sequence number of segment

Every segment (except for the first one) includes a sequence number so that the loss of a segment could be detected in reassembly. The sequence numbering shall start from zero with every new CPCS SDU. The first segment which contains an SAR.SEQ field must have SAR.SEQ = 0. All subsequent segments from the same CPCS SDU shall increase this sequence number so that the SAR.SEQ field adds one with every transmission.

The value SAR.NSEGS indicates the total number of segments, minus one. So when SAR.NSEGS = 0, the CPCS SDU is sent in one segment. SAR.NSEGS = 63 indicates there will be 64 segments to form the full CPCS SDU. When SAR.NSEGS = 0, it indicates that this first segment is also the last segment. No further segment with SAR.TYPE = 0b01 or 0b10 is to be expected for this one-segment CPCS SDU.

SAR at the receiving end shall buffer all segments and deliver only fully reassembled CPCS SDUs to the SSCS above. Should reassembly fail due to a segment not being received or too many segments being received, etc., SAR shall not deliver any incomplete CPCS SDU to the SSCS above.

9.2.2.2 SAR constants

Table 9-2 shows the constants for the SAR service.

Table 9-2 – SAR constants

Constant	Value
CIMTUSize	256 bytes
CIMaxAppPktSize	Max value (SAR.NSEGS) x CIMTUSize

9.3 NULL specific service convergence sublayer (NULL SSCS)

9.3.1 Overview

Null SSCS provides the MAC layer with a transparent path to upper layers, being as simple as possible and minimizing overheads. It is intended for applications that do not need any special convergence capability.

The unicast and multicast connections of this SSCS shall use the SAR service, as defined in clause 9.2.2. If they do not need the SAR service, they shall still include the SAR header (notifying just one segment).

The CON.TYPE and MUL.TYPE (see Annex C) for unicast connections and multicast groups shall use the same type that has already been defined for the application that makes use of this null SSCS.

9.3.2 Primitives

Null SSCS primitives are just a direct mapping of the MAC primitives. A full description of every primitive is avoided, because the mapping is direct and they will work as the ones of the MAC layer.

The directly mapped primitives have exactly the same parameters as the ones in the MAC layer and perform the same functionality. The set of primitives that are directly mapped are shown below.

Table 9-3 – Primitive mapping between the null SSCS primitives and the MAC layer primitives

Null SSCS mapped to a MAC primitive
CL_NULL_ESTABLISH.request	MAC_ESTABLISH.request
CL_NULL_ESTABLISH.indication	MAC_ESTABLISH.indication
CL_NULL_ESTABLISH.response	MAC_ESTABLISH.response
CL_NULL_ESTABLISH.confirm	MAC_ESTABLISH.confirm
CL_NULL_RELEASE.request	MAC_RELEASE.request
CL_NULL_RELEASE.indication	MAC_RELEASE.indication
CL_NULL_RELEASE.response	MAC_RELEASE.response
CL_NULL_RELEASE.confirm	MAC_RELEASE.confirm
CL_NULL_JOIN.request	MAC_JOIN.request
CL_NULL_JOIN.indication	MAC_JOIN.indication
CL_NULL_JOIN.response	MAC_JOIN.response
CL_NULL_JOIN.confirm	MAC_JOIN.confirm
CL_NULL_LEAVE.request	MAC_LEAVE.request
CL_NULL_LEAVE.indication	MAC_LEAVE.indication
CL_NULL_LEAVE.response	MAC_LEAVE.response
CL_NULL_LEAVE.confirm	MAC_LEAVE.confirm
CL_NULL_DATA.request	MAC_DATA.request
CL_NULL_DATA.indication	MAC_DATA.indication
CL_NULL_DATA.confirm	MAC_DATA.confirm
CL_NULL_SEND.request	MAC_SEND.request
CL_NULL_SEND.indication	MAC_SEND.indication
CL_NULL_SEND.confirm	MAC_SEND.confirm

9.4 IPv4 specific service convergence sublayer (IPv4 SSCS)

9.4.1 Overview

The IPv4 SSCS provides an efficient method for transferring IPv4 packets over the ITU-T G.9904 subnetworks. Several conventions do apply:

- A service node can send IPv4 packets to the base node or to other service nodes.
- It is assumed that the base node acts as a router between the ITU-T G.9904 subnetwork and any other network. The base node could also act as a NAT. How the base node connects to the other networks is beyond the scope of this Recommendation.
- In order to keep implementations simple, only one single route is supported per local IPv4 address.
- Service Nodes may use statically configured IPv4 addresses or DHCP to obtain IPv4 addresses.
- The base node performs IPv4 to EUI-48 address resolution. Each service node registers its IPv4 address and EUI-48 address with the base node (see clause 9.4.2). Other service nodes can then query the base node to resolve an IPv4 address into an EUI-48 address. This requires the establishment of a dedicated connection with the base node for address resolution.
- The IPv4 SSCS performs the routing of IPv4 packets. In other words, the IPv4 SSCS will decide whether the packet should be sent directly to another service node or forwarded to the configured gateway.
- Although IPv4 is a connectionless protocol, the IPv4 SSCS is connection-oriented. Once address resolution has been performed, a connection is established between the source and destination service node for the transfer of IPv4 packets. This connection is maintained while traffic is being transferred and may be closed after a period of inactivity.
- The CPCS (see clause 9.2) SAR sublayer shall always be present with the IPv4 convergence layer. Generated MSDUs are at most 'CIMTUSize' bytes long and upper layer PDU messages are not expected to be longer than CIMaxAppPktSize.
- Optionally TCP/IPv4 headers may be compressed. Compression is negotiated as part of the connection establishment phase.
- The broadcasting of IPv4 packets is supported using the MAC broadcast mechanism.
- The multicasting of IPv4 packets is supported using the MAC multicast mechanism.

The IPv4 SSCS has a number of connection types. For address resolution there is a connection to the base node. For IPv4 data transfer there is one connection per destination node: with the base node that acts as the IPv4 gateway to other networks or to/with any other node in the same subnetwork. This is shown in Figure 9-3.

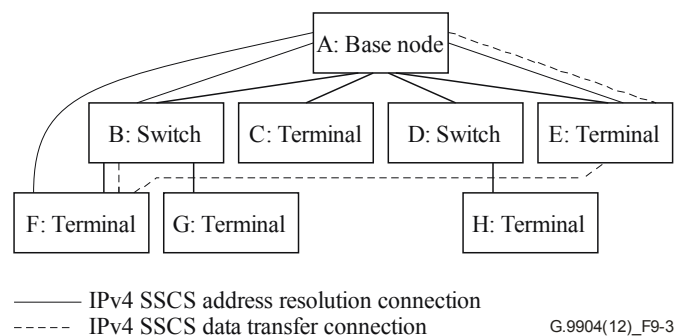


Figure 9-3 – IPv4 SSCS connection example

Here, nodes B, E and F have address resolution connections to the base node. Node E has a data connection to the base node and node F. Node F also has a data connection to node B. The figure does not show broadcast and multicast connections.

9.4.2 Address resolution

9.4.2.1 General

The IPv4 layer will present the IPV4 SSCS with an IPv4 packet to be transferred. The IPV4 SSCS is responsible for determining which service node the packet should be delivered to using the IPv4 addresses in the packet. The IPV4 SSCS must then establish a connection to the destination if one does not already exist so that the packet can be transferred. Three classes of IPv4 addresses can be used and the following subclauses describe how these addresses are resolved into EUI-48 addresses.

9.4.2.2 Unicast addresses

9.4.2.2.1 General

IPv4 unicast addresses must be resolved into unicast EUI-48 addresses. The base node maintains a database of IPv4 addresses and EUI-48 addresses. Address resolution then operates by querying this database. A service node must establish a connection to the address resolution service running on the base node, using the connection type value TYPE (see Annex C) TYPE_CL_IPv4_AR. No data should be passed in the connection establishment. Using this connection, the service node can use two mechanisms as defined in the following paragraphs.

9.4.2.2.2 Address registration and unregistration

A service node uses the AR_REGISTER_S message to register an IPv4 address and the corresponding EUI-48 address; this means a request from the base node to record inside its registration table, the IPv4 address and its corresponding service node EUI-48. The base node will acknowledge an AR_REGISTER_B message. The service node may register multiple IPv4 addresses for the same EUI-48 address.

A service node uses the AR_DEREGISTER_S message to unregister an IPv4 address and the corresponding EUI-48 address; this means requests from the base node to delete inside its registration table, the entry corresponding to the concerned IPv4 address. The base node will acknowledge it with an AR_DEREGISTER_B message.

When the IPv4 address resolution connection between the service node and the base node is closed, the base node should remove all addresses associated with that connection.

9.4.2.2.3 Address look-up

A service node uses the AR_LOOKUP_S message to perform a look-up. The message contains the IPv4 address to be resolved. The base node will respond with an AR_LOOKUP_B message that contains an error code and, if there is no error, the EUI-48 address associated with the IPv4 address. If the base node has multiple entries in its database for the same IPv4 address, the possible returned EUI-48 address is undefined.

9.4.2.3 Broadcast address

The IPv4 broadcast address 255.255.255.255 maps to a MAC broadcast connection with an LCID equal to LCI_CL_IPv4_BROADCAST. All IPv4 broadcast packets will be sent to this connection. When an IPv4 broadcast packet is received on this connection, the IPv4 address should be examined to determine if it is a broadcast packet for the subnetwork in which the node has an IPv4 address. Only broadcast packets from member subnets should be passed up the IPv4 protocol stack.

9.4.2.4 Multicast addresses

Multicast IPv4 addresses are mapped to an ITU-T G.9904 MAC multicast connection by the base node using an address resolution protocol.

To join a multicast group, AR_MCAST_REG_S is sent from the service node to the base node with the IPv4 multicast address. The base node will reply with an AR_MCAST_REG_B that contains the LCID value assigned to the said multicast address. However, the base node may also allocate other LCIDs which are not in use. The service node can then join a multicast group (see clause 8.6.7.2) for the given LCID to receive IPv4 multicast packets. These LCID values can be reused so that multiple IPv4 destination multicast addresses can be seen on the same LCID. To leave the multicast group, AR_MCAST_UNREG_S is sent from the service node to the base node with the IPv4 multicast address. The base node will acknowledge it with an AR_MCAST_UNREG_B message.

When a service node wants to send an IPv4 multicast datagram, it just uses the appropriate LCID. If the service node has not joined the multicast group, it first needs to learn the LCID to be used. The process with AR_MCAST_REG_{S|B} messages as described above can be used. While IPv4 multicast packets are still being sent, the service node remains registered to the multicast group. $T_{\text{mcast_reg}}$ after the last IPv4 multicast datagram was sent, the service node should unregister from the multicast group, by means of AR_MCAST_UNREG_{S|B} messages. The nominal value of $T_{\text{mcast_reg}}$ is 10 minutes; however, other values may be used.

9.4.2.5 Retransmission of address resolution packets

The connection between the service node and the base node for address resolution is not reliable if the MAC ARQ is not used. The service node is responsible for making retransmissions if the base node does not respond in one second. It is not considered an error when the base node receives the same registration requests multiple times or is asked to remove a registration that does not exist. These conditions can be the result of retransmissions.

9.4.3 IPv4 packet transfer

For packets to be transferred, a connection needs to be established between the source and destination nodes. The IPV4 SSCS will examine each IPv4 packet to determine the destination EUI-48 address. If a data connection to the destination already exists, the packet is sent. To establish this, IPV4 SSCS keeps a table for each connection, with the information shown in Table 9-4 (see [IETF RFC 1144]). To use this table, it is first necessary to determine if the IPv4 destination address is in the local subnetwork or if a gateway has to be used. The netmask associated with the local IPv4 address is used to determine this. If the IPv4 destination address is not in the local subnetwork, the address of the default gateway is used instead of the destination address when the table is searched.

Table 9-4 – IPV4 SSCS table entry

Parameter	Description
CL_IPv4_Con.Remote_IP	Remote IPv4 address of this connection
CL_IPv4_Con.ConHandle	MAC connection handle for the connection
CL_IPv4_Con.LastUsed	Timestamp of last packet received/transmitted
CL_IPv4_Con.HC	Header compression scheme being used
CL_IPv4_CON.RxSeq	Next expected receive sequence number
CL_IPv4_CON.TxSeq	Sequence number for next transmission

The IPV4 SSCS may close a connection when it has not been used for an implementation-defined time period. When the connection is closed the entry for the connection is removed at both ends of the connection.

When a connection to the destination does not exist, more work is necessary. The address resolution service is used to determine the EUI-48 address of the remote IPv4 address if it is local or the gateway associated with the local address if the destination address is in another subnetwork. When the base node replies with the EUI-48 address of the destination service node, a MAC connection is established with the remote device. The TYPE value of this connection is TYPE_CL_IPv4_UNICAST. The data passed in the request message is defined in clause 9.4.7.4. The local IPv4 address is provided so that the remote device can add the new connection to its cache of connections for sending data in the opposite direction. The use of Van Jacobson (VJ) header compression is also negotiated as part of the connection establishment. Once the connection has been established, the IPv4 packet can be sent.

When the packet is addressed to the IPv4 broadcast address, the packet has to be sent using the MAC broadcast service. When the IPV4 SSCS is opened, a broadcast connection is established for transferring all broadcast packets. The broadcast IPv4 packet is simply sent to this connection. Any packet received on this broadcast connection is passed to the IPv4 protocol stack.

9.4.4 Segmentation and reassembly

The IPV4 SSCS should support IPv4 packets with an MTU of 1500 bytes. This requires the use of SAR (see clause 9.2.2).

9.4.5 Header compression

Van Jacobson TCP/IP header compression is an optional feature in the IPv4 SSCS. The use of VJ compression is negotiated as part of the connection establishment phase of the connection between two service nodes.

VJ compression is designed for use over a point-to-point link layer that can inform the decompressor when packets have been corrupted or lost. When there are errors or lost packets, the decompressor can then resynchronize with the compressor. Without this resynchronization process, erroneous packets will be produced and passed up the IPv4 stack.

The MAC layer does not provide the facility of detecting lost packets or reporting corrupt packets. Thus, it is necessary to add this functionality in the IPV4 SSCS. The IPV4 SSCS maintains two sequence numbers when VJ compression is enabled for a connection. These sequence numbers are 8 bits in size. When transmitting an IPv4 packet, the CL_IPv4_CON.TxSeq sequence number is placed in the packet header, as shown in clause 9.4.3. The sequence number is then incremented. Upon receipt of a packet, the sequence number in the received packet is compared against CL_IPv4_CON.RxSeq. If they differ, TYPE_ERROR, as defined in [IETF RFC1144], is passed to the decompressor. The CL_IPv4_CON.RxSeq value is always updated to the value received in the packet header.

Header compression should never be negotiated for broadcast or multicast packets.

9.4.6 Quality of service mapping

The ITU-T G.9904 MAC specifies that the contention-based access mechanism supports four priority levels (1-4). Level 1 is used for MAC signalling messages, but not exclusively so.

IPv4 packets include a type of service (TOS) field in the header to indicate the QoS the packet would like to receive. Three bits of the TOS indicate the IP precedence. The following table specifies how the IP precedence is mapped into the ITU-T G.9904 MAC priority.

Table 9-5 – Mapping IPv4 precedence to ITU-T G.9904 MAC priority

IP precedence	MAC priority
000 – Routine	4
001 – Priority	4
010 – Immediate	3
011 – Flash	3
100 – Flash Override	2
101 – Critical	2
110 – Internetwork Control	1
111 – Network Control	1

9.4.7 Packet formats and connection data

9.4.7.1 General

This clause defines the format of IPV4 SSCS PDUs.

9.4.7.2 Address resolution PDUs

9.4.7.2.1 General

The following PDUs are transferred over the address resolution connection between the service node and the base node. The following clauses define AR.MSG values in the range of 0 to 11. All higher values are reserved for later versions of this Recommendation.

9.4.7.2.2 AR_REGISTER_S

Table 9-6 shows the address resolution register message sent from the service node to the base node.

Table 9-6 – AR_REGISTER_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type • For AR_REGISTER_S = 0
AR.IPv4	32-bits	IPv4 address to be registered
AR.EUI-48	48-bits	EUI-48 to be registered

9.4.7.2.3 AR_REGISTER_B

Table 9-7 shows the address resolution register acknowledgment message sent from the base node to the service node.

Table 9-7 – AR_REGISTER_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type • For AR_REGISTER_B = 1
AR.IPv4	32-bits	Registered IPv4 address
AR.EUI-48	48-bits	EUI-48 registered

The AR.IPv4 and AR.EUI-48 fields are included in the AR_REGISTER_B message so that the service node can perform multiple overlapping registrations.

9.4.7.2.4 AR_UNREGISTER_S

Table 9-8 shows the address resolution unregister message sent from the service node to the base node.

Table 9-8 – AR_UNREGISTER_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type. • For AR_UNREGISTER_S = 2
AR.IPv4	32-bits	IPv4 address to be unregistered
AR.EUI-48	48-bits	EUI-48 to be unregistered

9.4.7.2.5 AR_UNREGISTER_B

Table 9-9 shows the address resolution unregister acknowledgment message sent from the base node to the service node.

Table 9-9 – AR_UNREGISTER_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type • For AR_UNREGISTER_B = 3
AR.IPv4	32-bits	Unregistered IPv4 address
AR.EUI-48	48-bits	Unregistered EUI-48

The AR.IPv4 and AR.EUI-48 fields are included in the AR_UNREGISTER_B message so that the service node can perform multiple overlapping unregistrations.

9.4.7.2.6 AR_LOOKUP_S

Table 9-10 shows the address resolution look-up message sent from the service node to the base node.

Table 9-10 – AR_LOOKUP_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type • For AR_LOOKUP_S = 4
AR.IPv4	32-bits	IPv4 address to look up

9.4.7.2.7 AR_LOOKUP_B

Table 9-11 shows the address resolution look-up response message sent from the base node to the service node.

Table 9-11 – AR_LOOKUP_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none"> • For AR_LOOKUP_B = 5
AR.IPv4	32-bits	IPv4 address looked up
AR.EUI-48	48-bits	EUI-48 for IPv4 address
AR.Status	8-bits	Look-up status, indicating if the address was found or an error occurred: <ul style="list-style-type: none"> • 0 = found, AR.EUI-48 valid • 1 = unknown, AR.EUI-48 undefined.

The look-up may fail if the requested address has not been registered. In that case, AR.Status will have a value other than zero and the contents of AR.EUI-48 will be undefined. The look-up is only successful when AR.Status is zero. In that case, the EUI-48 field contains the resolved address.

9.4.7.2.8 AR_MCAST_REG_S

Table 9-12 shows the multicast address resolution register message sent from the service node to the base node.

Table 9-12 – AR_MCAST_REG_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none"> • For AR_MCAST_REG_S = 8
AR.IPv4	32-bits	IPv4 multicast address to be registered

9.4.7.2.9 AR_MCAST_REG_B

Table 9-13 shows the multicast address resolution register acknowledgment message sent from the base node to the service node.

Table 9-13 – AR_MCAST_REG_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none"> • For AR_MCAST_REG_B = 9
AR.IPv4	32-bits	IPv4 multicast address registered
Reserved	2-bits	Reserved by ITU-T. Should be encoded as 0
AR.LCID	6-bits	LCID assigned to this IPv4 multicast address

The AR.IPv4 field is included in the AR_MCAST_REG_B message so that the service node can perform multiple overlapping registrations.

9.4.7.2.10 AR_MCAST_UNREG_S

Table 9-14 shows the multicast address resolution unregister message sent from the service node to the base node.

Table 9-14 – AR_MCAST_UNREG_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none"> For AR_MCAST_UNREG_S = 10
AR.IPv4	32-bits	IPv4 multicast address to be unregistered

9.4.7.2.11 AR_MCAST_UNREG_B

Table 9-15 shows the multicast address resolution unregister acknowledgment message sent from the base node to the service node.

Table 9-15 – AR_MCAST_UNREG_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none"> For AR_MCAST_UNREG_B = 11
AR.IPv4	32-bits	IPv4 multicast address unregistered

The AR.IPv4 field is included in the AR_MCAST_UNREG_B message so that the service node can perform multiple overlapping unregistrations.

9.4.7.3 IPv4 packet format**9.4.7.3.1 General**

The following PDU formats are used for transferring IPv4 packets between service nodes. Two formats are defined. The first format is for when header compression is not used. The second format is for Van Jacobson header compression.

9.4.7.3.2 IPv4 packet format, no negotiated header compression

When no header compression has been negotiated, the IPv4 packet is simply sent as is, without any header.

Table 9-16 – IPv4 packet format without negotiated header compression

Name	Length	Description
IPv4.PKT	n-octets	The IPv4 packet

9.4.7.3.3 IPv4 packet format with VJ header compression

With Van Jacobsen header compression, a one-octet header is needed before the IPv4 packet.

Table 9-17 – IPv4 packet format with VJ header compression negotiated

Name	Length	Description
IPv4.Type	2-bits	Type of compressed packet: <ul style="list-style-type: none"> IPv4.Type = 0 – TYPE_IP IPv4.Type = 1 – UNCOMPRESSED_TCP IPv4.Type = 2 – COMPRESSED_TCP IPv4.Type = 3 – TYPE_ERROR.
IPv4.Seq	6-bits	Packet sequence number
<i>IPv4.PKT</i>	n-octets	The IPv4 packet

The IPv4.Type value TYPE_ERROR is never sent. It is a pseudo packet type used to tell the decompressor that a packet has been lost.

9.4.7.4 Connection data

9.4.7.4.1 General

When a connection is established between service nodes for the transfer of IPv4 packets, data is also transferred in the connection request packets. This data allows the negotiation of compression and notification of the IPv4 address.

9.4.7.4.2 Connection data from the initiator

Table 9-18 shows the connection data sent by the initiator.

Table 9-18 – Connection signalling data sent by the initiator

Name	Length	Description
Reserved	6-bits	Reserved by ITU-T. Should be encoded as 0 in this version of the IPV4 SSCS protocol.
Data.HC	2-bit	Header compression <ul style="list-style-type: none"> • Data.HC = 0 – No compression requested • Data.HC = 1 – VJ Compression requested • Data.HC = 2, 3 – Reserved by ITU-T.
Data.IPv4	32-bits	IPv4 address of the initiator

If the device accepts the connection, it should copy the Data.IPv4 address into a new table entry along with the negotiated Data.HC value.

9.4.7.4.3 Connection data from the responder

Table 9-19 shows the connection data sent in response to the connection request.

Table 9-19 – Connection signalling data sent by the responder

Name	Length	Description
Reserved	6-bits	Should be encoded as zero in this version of the IPV4 SSCS protocol.
Data.HC	2-bit	Header compression negotiated: <ul style="list-style-type: none"> • Data.HC = 0 – No compression permitted • Data.HC = 1 – VJ Compression negotiated • Data.HC = 2,3 – Reserved by ITU-T.

A header compression scheme can only be used when it is supported by both service nodes. The responder may only set Data.HC to 0 or the same value as that received from the initiator. When the same value is used, it indicates that the requested compression scheme has been negotiated and will be used for the connection. Setting Data.HC to 0 allows the responder to deny the request for that header compression scheme or force the use of no header compression.

9.4.8 Service access point

9.4.8.1 General

This clause defines the service access point used by the IPv4 layer to communicate with the IPV4 SSCS.

9.4.8.2 Opening and closing the IPv4 SSCS

9.4.8.2.1 General

The following primitives are used to open and close the IPv4 SSCS. The IPv4 SSCS may be opened once only. The IPv4 layer may close the IPv4 SSCS when the IPv4 interface is brought down. The IPv4 SSCS will also close the IPv4 SSCS when the underlying MAC connection to the base node has been lost.

9.4.8.2.2 CL_IPv4_ESTABLISH.request

The CL_IPv4_ESTABLISH.request primitive is passed from the IPv4 layer to the IPV4 SSCS. It is used when the IPv4 layer brings the interface up.

The semantics of this primitive are as follows:

CL_IPv4_ESTABLISH.request{}

On receiving this primitive, the IPV4 SSCS will form the address resolution connection to the base node and join the broadcast group used for receiving/transmitting broadcast packets.

9.4.8.2.3 CL_IPv4_ESTABLISH.confirm

The CL_IPv4_ESTABLISH.confirm primitive is passed from the IPV4 SSCS to the IPv4 layer. It is used to indicate that the IPv4 SSCS is ready to access IPv4 packets to be sent to peers.

The semantics of this primitive are as follows:

CL_IPv4_ESTABLISH.confirm{}

Once the IPv4 SSCS has established all the necessary connections and is ready to transmit and receive IPv4 packets, this primitive is passed to the IPv4 layer. If the IPV4 SSCS encounters an error while opening, it responds with a CL_IPv4_RELEASE.confirm primitive, rather than a CL_IPv4_ESTABLISH.confirm.

9.4.8.2.4 CL_IPv4_RELEASE.request

The CL_IPv4_RELEASE.request primitive is used by the IPv4 layer when the interface is put down. The IPV4 SSCS closes all connections so that no more IPv4 packets are received and all resources are released.

The semantics of this primitive are as follows:

CL_IPv4_RELEASE.request{}

Once the IPV4 SSCS has released all its connections and resources it returns a CL_IPv4_RELEASE.confirm.

9.4.8.2.5 CL_IPv4_RELEASE.confirm

The CL_IPv4_RELEASE.confirm primitive is used by the IPv4 SSCS to indicate to the IPv4 layer that the IPv4 SSCS has been closed. This can be as a result of a CL_IPv4_RELEASE.request primitive, a CL_IPv4_ESTABLISH.request primitive, or because the MAC layer indicates the address resolution connection has been lost, or the service node itself is no longer registered.

The semantics of this primitive are as follows:

CL_IPv4_RELEASE.confirm{Result}

The result parameter has the meanings defined in Table C.3.

9.4.8.3 Unicast address management

9.4.8.3.1 General

The primitives defined here are used for address management, i.e., the registration and unregistration of IPv4 addresses associated with this IPv4 SSCS.

When there are no IPv4 addresses associated with the IPv4 SSCS, the IPv4 SSCS will only send and receive broadcast and multicast packets; unicast packets may not be sent. However, this is sufficient for BOOTP/DHCP operation to allow the device to gain an IPv4 address. Once an IPv4 address has been registered, the IPv4 layer can transmit unicast packets that have a source address equal to one of its registered addresses.

9.4.8.3.2 CL_IPv4_REGISTER.request

This primitive is passed from the IPv4 layer to the IPv4 SSCS to register an IPv4 address.

The semantics of this primitive are as follows:

CL_IPv4_REGISTER.request{IPv4, netmask, gateway}

The IPv4 address is the address to be registered.

The netmask is the network mask, used to mask the network number from the address. The netmask is used by the IPv4 SSCS to determine whether the packet should be delivered directly or the gateway should be used.

The gateway is an IPv4 address of the gateway to be used for packets with the IPv4 local address but the destination address is not in the same subnetwork as the local address.

Once the IPv4 address has been registered to the base node, a CL_IPv4_REGISTER.confirm primitive is used. If the registration fails, the CL_IPv4_RELEASE.confirm primitive will be used.

9.4.8.3.3 CL_IPv4_REGISTER.confirm

This primitive is passed from the IPv4 SSCS to the IPv4 layer to indicate that a registration has been successful.

The semantics of this primitive are as follows:

CL_IPv4_REGISTER.confirm{IPv4}

The IPv4 address is the address that was registered.

Once registration has been completed, the IPv4 layer may send IPv4 packets using this source address.

9.4.8.3.4 CL_IPv4_UNREGISTER.request

This primitive is passed from the IPv4 layer to the IPv4 SSCS to unregister an IPv4 address.

The semantics of this primitive are as follows:

CL_IPv4_UNREGISTER.request{IPv4}

The IPv4 address is the address to be unregistered.

Once the IPv4 address has been unregistered to the base node, a CL_IPv4_UNREGISTER.confirm primitive is used. If the unregistration fails, the CL_IPv4_RELEASE.confirm primitive will be used.

9.4.8.3.5 CL_IPv4_UNREGISTER.confirm

This primitive is passed from the IPv4 SSCS to the IPv4 layer to indicate that an unregistration has been successful.

The semantics of this primitive are as follows:

CL_IPv4_UNREGISTER.confirm{IPv4}

The IPv4 address is the address that was unregistered.

Once unregistration has been completed, the IPv4 layer may not send IPv4 packets using this source address.

9.4.8.4 Multicast group management

9.4.8.4.1 General

This clause describes the primitives used to manage multicast groups.

9.4.8.4.2 CL_IPv4_IGMP_JOIN.request

This primitive is passed from the IPv4 layer to the IPv4 SSCS. It contains an IPv4 multicast address that is to be joined.

The semantics of this primitive are as follows:

CL_IPv4_IGMP_JOIN.request{IPv4 }

The IPv4 address is the IPv4 multicast group that is to be joined.

When the IPv4 SSCS receives this primitive, it will arrange for IPv4 packets sent to this group to be multicast in the ITU-T G.9904 network and receive packets using this address to be passed to the IPv4 stack. If the IPv4 SSCS cannot join the group, it uses the CL_IPv4_IGMP_LEAVE.confirm primitive. Otherwise the CL_IPv4_IGMP_JOIN.confirm primitive is used to indicate success.

9.4.8.4.3 CL_IPv4_IGMP_JOIN.confirm

This primitive is passed from the IPv4 SSCS to the IPv4. It contains a result status and an IPv4 multicast address that was joined.

The semantics of this primitive are as follows:

CL_IPv4_IGMP_JOIN.confirm{IPv4}

The IPv4 address is the IPv4 multicast group that was joined. The IPv4 SSCS will start forwarding IPv4 multicast packets for the given multicast group.

9.4.8.4.4 CL_IPv4_IGMP_LEAVE.request

This primitive is passed from the IPv4 layer to the IPv4 SSCS. It contains an IPv4 multicast address to be left.

The semantics of this primitive are as follows:

CL_IPv4_IGMP_LEAVE.request{IPv4}

The IPv4 address is the IPv4 multicast group to be left. The IPv4 SSCS will stop forwarding IPv4 multicast packets for this group and may leave the ITU-T G.9904 MAC multicast group.

9.4.8.4.5 CL_IPv4_IGMP_LEAVE.confirm

This primitive is passed from the IPv4 SSCS to the IPv4. It contains a result status and an IPv4 multicast address that was left.

The semantics of this primitive are as follows:

CL_IPv4_IGMP_LEAVE.confirm{IPv4, Result}

The IPv4 address is the IPv4 multicast group that was left. The IPv4 SSSC will stop forwarding IPv4 multicast packets for the given multicast group.

The result takes a value from Table C.3.

This primitive can be used by the IPv4 SSSC as a result of a *CL_IPv4_IGMP_JOIN.request*, *CL_IPv4_IGMP_LEAVE.request* or because of an error condition resulting in the loss of the ITU-T G.9904 MAC multicast connection.

9.4.8.5 Data transfer

9.4.8.5.1 General

The following primitives are used to send and receive IPv4 packets.

9.4.8.5.2 CL_IPv4_DATA.request

This primitive is passed from the IPv4 layer to the IPv4 SSSC. It contains one IPv4 packet to be sent.

The semantics of this primitive are as follows:

CL_IPv4_DATA.request{IPv4_PDU}

The IPv4_PDU is the IPv4 packet to be sent.

9.4.8.5.3 CL_IPv4_DATA.confirm

This primitive is passed from the IPv4 SSSC to the IPv4 layer. It contains a status indication and an IPv4 packet that has just been sent.

The semantics of this primitive are as follows:

CL_IPv4_DATA.confirm{IPv4_PDU, Result}

The IPv4_PDU is the IPv4 packet that was to be sent.

The result value indicates whether the packet was sent or an error occurred. It takes a value from Table C.3.

9.4.8.5.4 CL_IPv4_DATA.indicate

This primitive is passed from the IPv4 SSSC to the IPv4 layer. It contains an IPv4 packet that has just been received.

The semantics of this primitive are as follows:

CL_IPv4_DATA.indicate{IPv4_PDU}

The IPv4_PDU is the IPv4 packet that was received.

9.5 IEC 61334-4-32 specific service convergence sublayer (IEC 61334-4-32 SSSC)

9.5.1 General

For all the service required, the IEC 61334-4-32 SSSC supports the DL_DATA primitives as defined in the IEC 61334-4-32 standard. IEC 61334-4-32 should be read at the same time as this clause, which is not standalone text.

9.5.2 Overview

The IEC 61334-4-32 SSCS provides convergence functions for applications that use IEC 61334-4-32 services. Implementations conforming to this SSCS shall offer all LLC basic and management services as specified in [IEC 61334-4-32], clauses 2.2.1 and 2.2.3. Additionally, the IEC 61334-4-32 SSCS specified in this clause provides extra services that help mapping this connection-less IEC 61334-4-32 LLC protocol to the connection-oriented nature of the MAC.

- A service node can only exchange data with the base node and not with other service nodes. This meets all the requirements of [IEC 61334-4-32], which has similar restrictions.
- Each IEC 61334-4-32 SSCS session establishes a dedicated ITU-T G.9904 MAC connection for exchanging unicast data with the base node.
- The service node SSCS session is responsible for initiating this connection to the base node. The base node SSCS cannot initiate a connection to a service node.
- Each IEC 61334-4-32 SSCS listens to an ITU-T G.9904 broadcast MAC connection dedicated to the transfer of IEC 61334-4-32 broadcast data from the base node to the service nodes. This broadcast connection is used when applications in the base node using IEC 61334-4-32 services make a transmission request with the Destination_address used for broadcast or the broadcast SAP functions are used. When there are multiple SSCS sessions within a service node, one ITU-T G.9904 broadcast MAC connection is shared by all the SSCS sessions.
- A CPCS session is always present with an IEC 61334-4-32 SSCS session. The SPCS sublayer functionality is as specified in clause 9.2.2. Thus, the MSDUs generated by IEC 61334-4-32 SSCS are always less than CIMTUSize bytes and application messages shall not be longer than CIMaxAppPktSize.

9.5.3 Address allocation and connection establishment

Each 4-32 connection will be identified with the "Application unique identifier" that will be communicating through this 4-32 connection. It is the scope of the communication profile based on these lower layers to define the nature and rules for, this unique identifier. As long as the specification of the 4-32 convergence layer concerns this identifier, it will be called the "Device Identifier".

The protocol stack as defined in IEC 61334 defines a destination address to identify each device in the network. This destination address is specified beyond the scope of the IEC 61334-4-32 document. However, it is used by the document. So that ITU-T G.9904 devices can make use of the 4-32 layer, this destination address is also required and is specified here. For more information about this destination address, please see [IEC 61334-4-1] clause 4.3, MAC addresses.

The destination address has a scope of one ITU-T G.9904 subnetwork. The base node 4-32 SSCP layer is responsible for allocating these addresses dynamically and associating the device identifier of the service nodes SSCP session device with the allocated destination address, according to the IEC 61334-4-1 standard. The procedure is as follows:

When the service node IEC 61334-4-32 SSCS session is opened by the application layer, it passes the device identifier of the device. The IEC 61334-4-32 SSCS session then establishes its unicast connection to the base node. This unicast connection uses the ITU-T G.9904 MAC TYPE value TYPE_CL_432, as defined in Table C.1. The connection request packet sent from the service node to the base node contains a data parameter. This data parameter contains the device identifier. The format of this data is specified in clause 9.5.4.2.

On receiving this connection request at the base node, the base node allocates a unique subnetwork destination address to the service nodes SSCS session. The base node sends back an ITU-T G.9904 MAC connection response packet that contains a data parameter. This data parameter contains the allocated destination address and the address being used by the base node itself. The format of this

data parameter is defined in clause 9.5.4.2. A 4-32 CL SAP primitive is used in the base node to indicate this new service node SSCS session mapping of the device identifier and Destination_address to the 4-32 application running in the base node.

On receiving the connection establishment and the Destination_address passed in the ITU-T G.9904 MAC connection establishment packet, the 4-32 SSCS session confirms to the application that the convergence layer session has been opened and indicates the Destination_address allocated to the service node SSCS session and the address of the base node. The service node also opens an ITU-T G.9904 MAC broadcast connection with the LCID equal to LCI_CL_432_BROADCAST, as defined in Table C.2, if no other SSCS session has already opened such a broadcast connection. This connection is used to receive broadcast packets sent by the base node 4-32 convergence layer to all service node 4-32 convergence layer sessions.

If the base node has allocated all its available Destination_addresses, due to the exhaustion of the address space or implementation limits, it should simply reject the connection request from the service node. The service node may try to establish the connection again. However, to avoid overloading the ITU-T G.9904 subnetwork with such requests, it should limit such connection establishments to one attempt per minute when the base node rejects a connection establishment.

When the unicast connection between a service node and the base node is closed (e.g., because the convergence layer on the service node is closed or the ITU-T G.9904 MAC level connection between the service node and the base node is lost), the base node will de-allocate the Destination_address allocated to the service node SSCS session. The base node will use a 4-32 CL SAP (CL_432_Leave.indication) primitive to indicate the de-allocation of the Destination_address to the 4-32 application running on the base node.

9.5.4 Connection establishment data format

9.5.4.1 General

As described in clause 9.5.3, the MAC ITU-T G.9904 connection data is used to transfer the device identifier to the base node and the allocated Destination_address to the service node SSCS session. This clause describes the format used for this data.

9.5.4.2 Service node to base node

The service node session passes the device identifier to the base node as part of the connection establishment request. The format of this message is shown in Table 9-20.

Table 9-20 – Connection signalling data sent by the service node

Name	Length	Description
Data.SN	n-Octets	Device identifier "COSEM logical device name" of the "Management logical device" of the DLMS/COSEM device as specified in the DLMS/COSEM, which will be communicating through this 4-32 connection.

9.5.4.3 Base node to service node

The base node passes the allocated Destination_address to the service node session as part of the connection establishment request. It also gives its own address to the service node. The format of this message is shown in Table 9-21.

Table 9-21 – Connection signalling data sent by the base node

Name	Length	Description
<i>Reserved</i>	4-bits	Reserved by ITU-T. Should be encoded as zero in this version of the Recommendation.
Data.DA	12-bits	Destination_address allocated to the service node
<i>Reserved</i>	4-bits	Reserved by ITU-T. Should be encoded as zero in this version of the Recommendation.
Data.BA	12-bits	Base_address used by the base node

9.5.5 Packet format

The packet formats are used as defined in [IEC 61334-4-32], clause 4, LLC protocol data unit structure (LLC_PDU).

9.5.6 Service access point**9.5.6.1 Opening and closing the convergence layer at the service node****9.5.6.1.1 CL_432_ESTABLISH.request**

This primitive is passed from the application to the 4-32 convergence layer. It is used to open a convergence layer session and initiate the process of registering the device identifier with the base node and the base node allocating a Destination_address to the service node session.

The semantics of this primitive are as follows:

CL_432_ESTABLISH.request{ *DeviceIdentifier* }

The device identifier is that of the device to be registered with the base node.

If the device identifier is registered and the convergence layer session is successfully opened, the primitive CL_432_ESTABLISH.confirm is used. If an error occurs the primitive CL_432_RELEASE.confirm is used.

9.5.6.1.2 CL_432_ESTABLISH.confirm

This primitive is passed from the 4-32 convergence layer to the application. It is used to confirm the successful opening of the convergence layer session and that data may now be passed over the convergence layer.

The semantics of this primitive are as follows:

CL_432_ESTABLISH.confirm{ *DeviceIdentifier*, *Destination_address*, *Base_address* }

The device identifier is used to identify which CL_432_ESTABLISH.request this CL_432_ESTABLISH.confirm is for.

The Destination_address is the address allocated to the service node 4-32 session by the base node.

The Base_address is the address being used by the base node.

9.5.6.1.3 CL_432_RELEASE.request

This primitive is passed from the application to the 4-32 convergence layer. It is used to close the convergence layer and release any resources it may be holding.

The semantics of this primitive are as follows:

CL_432_RELEASE.request{*Destination_address*}

The Destination_address is the address allocated to the service node 4-32 session which is to be closed.

The convergence layer will use the primitive `CL_432_RELEASE.confirm` when the convergence layer session has been closed.

9.5.6.1.4 CL_432_RELEASE.confirm

This primitive is passed from the 4-32 convergence layer to the application. The primitive tells the application that the convergence layer session has been closed. This could be because of a `CL_432_RELEASE.request` or because an error has occurred, forcing the closure of the convergence layer session.

The semantics of this primitive are as follows:

CL_432_RELEASE.confirm{Destination_address, Result}

The handle identifies the session which has been closed.

The result parameter has the meanings defined in Table C.3.

9.5.6.2 Opening and closing the convergence layer at the base node

No service access point primitives are defined at the base node for opening or closing the convergence layer. None are required since the 4-32 application in the base node does not need to pass any information to the 4-32 convergence layer in the base node.

9.5.6.3 Base node indications

9.5.6.3.1 General

The following primitives are used in the base node 4-32 convergence layer to indicate events to the 4-32 application in the base node. They indicate when a service node session has joined or left the network.

9.5.6.3.2 CL_432_JOIN.indicate

CL_432_JOIN.indicate{ Device Identifier, Destination_address}

The device identifier is that of the device connected to the service node that has just joined the network.

The `Destination_address` is the address allocated to the service node by the base node.

9.5.6.3.3 CL_432_LEAVE.indicate

CL_432_LEAVE.indicate{Destination_address}

The `Destination_address` is the address of the service node session that has just left the network.

9.5.6.4 Data transfer primitives

The data transfer primitives are used as defined in [IEC 61334-4-32], clauses 2.2, 2.3, 2.4 and 2.11, LLC service specification. As stated earlier, ITU-T G.9904 SCS make the use of IEC 61334-4-32 DL_Data service (.req, .conf, .ind) for carrying out all the data involved during data transfer.

9.6 IPv6 service-specific convergence sublayer (IPv6 SCS)

9.6.1 Overview

9.6.1.1 General

The IPv6 convergence layer provides an efficient method for transferring IPv6 packets over the ITU-T G.9904 network.

A service node can pass IPv6 packets to the base node or directly to other service nodes.

By default, the base node acts as a router between the ITU-T G.9904 subnetwork and the backbone network. All the base nodes must have at least this connectivity capability. Any other node inside the subnetwork can also act as a gateway. The base node could also act as an NAT router. However given the abundance of IPv6 addresses this is not expected. How the base node connects to the backbone is beyond the scope of this standard.

9.6.1.2 IPv6 unicast addressing assignment

- IPv6 service nodes (and base nodes) shall support the standard IPv6 protocol, as described in [IETF RFC 2460].
- IPv6 service nodes (and base nodes) shall support the standard IPv6 addressing architecture, as described in [IETF RFC 4291].
- IPv6 service nodes (and base nodes) shall support global unicast IPv6 addresses, link-local IPv6 addresses and multicast IPv6 addresses, as described in [IETF RFC 4291].
- IPv6 service nodes (and base nodes) shall support automatic address configuration using stateless address configuration [IETF RFC 2462]. They may also support automatic address configuration using stateful address configuration [IETF RFC 3315] and they may support manual configuration of IPv6 addresses. The decision for which address configuration scheme to use is deployment specific.
- The service node shall support DHCPv6 client, when base nodes have to support a DHCPv6 server as described in [IETF RFC 3315] for stateless address configuration.

9.6.1.3 Address management in ITU-T G.9904 subnetworks

Packets are routed in ITU-T G.9904 subnetworks according to the node identifier NID. The node identifier is a combination of a service node's LNID and SID (see clause 8.2). The base node is responsible for assigning the LNID to service nodes. During the registration process which leads to LNID assignment to the related service node, the base node registers the service node EUI-48, and the assigned LNID together with the SID.

At the convergence layer level, addressing is performed using the EUI-48 of the related service node. The role of the convergence sublayer is to resolve the IPv6 address into EUI-48 of the service node. This is done using the address resolution service set of the base node.

9.6.1.4 Role of the base node

At the convergence sublayer level, the base node maintains a table containing all the IPv6 unicast addresses and the EUI-48 related to them. One of the roles of the base node is to perform IPv6 to EUI-48 address resolution. Each service node belonging to the subnetwork managed by the base node registers its IPv6 address and EUI-48 address with the base node. Other service nodes can then query the base node to resolve an IPv6 address into an EUI-48 address. This requires the establishment of a dedicated connection to the base node for address resolution, which is shared by both IPv4 and IPv6 address resolution.

Optionally UDP/IPv6 headers may be compressed. Compression is negotiated as part of the connection establishment phase. Currently, there is one header compression technique described in the present Recommendation that is used for the transmission of IPv6 packets over IEEE 802.15.4 networks, as defined in [SP 800-57]. This is also known as LOWPAN_IPHC1.

The multicasting of IPv6 packets is supported using the MAC multicast mechanism.

9.6.2 IPv6 convergence layer

9.6.2.1 Overview

9.6.2.1.1 General

The convergence layer has a number of connection types. For address resolution there is a connection to the base node. For IPv6 data transfer there is one connection per destination node: the base node that acts as the IPv6 gateway to the outside world or another node in the same subnetwork. This is shown in Figure 9-4.

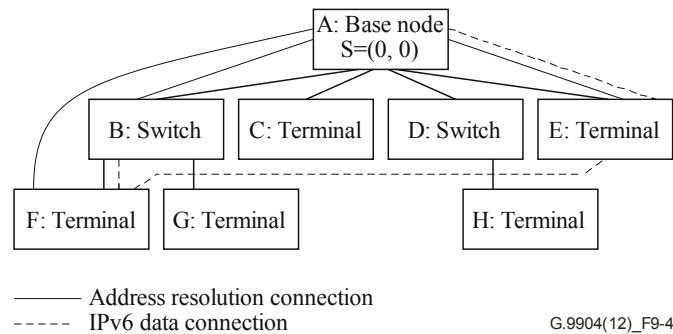


Figure 9-4 – Example of IPv6 connection

Here, nodes B, E and F have address resolution connections to the base node. Node E has a data connection to the base node and node F. Node F also has a data connection to node B. The figure does not show broadcast-traffic and multicast-traffic connections.

9.6.2.1.2 Routing in the subnetwork

Routing IPv6 packets is the scope of the convergence layer. In other words, the convergence layer will decide whether the packet should be sent directly to another service node or forwarded to the configured gateway depending on the IPv6 destination address.

Although IPv6 is a connectionless protocol, the IPv6 convergence layer is connection-oriented. Once address resolution has been performed, a connection is established between the source and destination service nodes for the transfer of IP packets. This connection is maintained all the time that the traffic is being transferred and may be removed after a period of inactivity.

9.6.2.1.3 SAR

The CPCS sublayer shall always be present with the IPv6 convergence layer allowing segmentation and reassembly facilities. The SAR sublayer functionality is given in clause 9.2. Thus, the MSDUs generated by the IPv6 convergence layer are always less than CIMITUSize bytes and application messages are expected to be no longer than CImaxAppPktSize.

9.6.3 IPv6 address configuration

9.6.3.1 Overview

The service nodes may use statically configured IPv6 addresses, link local addresses, stateless or stateful auto-configuration according to [IETF RFC 2462], or DHCPv6 to obtain IPv6 addresses. All the nodes shall support the unicast link local address, in addition to other configured addresses below, and multicast addresses, if ever the node belongs to multicast groups.

9.6.3.2 Interface identifier

In order to make use of stateless address auto configuration and link local addresses it is necessary to define how the interface identifier, as defined in [IETF RFC 4291], is derived. Each ITU-T G.9904 node has a unique EUI-48. This EUI-48 is converted into an EUI-64 in the same way as for Ethernet networks as defined in [IETF RFC 2464]. This EUI-64 is then used as the interface identifier.

9.6.3.3 IPv6 link local address configuration

The IPv6 link local address of an ITU-T G.9904 interface is formed by appending the interface identifier as defined above to the prefix FE80::/64.

9.6.3.4 Stateless address configuration

An IPv6 address prefix used for stateless auto configuration, as defined in [IETF RFC4862], of an ITU-T G.9904 interface shall have a length of 64 bits. The IPv6 prefix is obtained by the service nodes from the base node via router advertisement messages, which are sent periodically or on request by the base node.

9.6.3.5 Stateful address configuration

An IPv6 address can be alternatively configured using DHCPv6, as described in [IETF RFC 3315]. DHCPv6 can provide a device with addresses assigned by a DHCPv6 server and other configuration information, which are carried in options.

9.6.3.6 Multicast address

IPv6 service nodes (and base nodes) shall support the multicast IPv6 addressing, as described in [IETF RFC 4291], clause 2.7.

9.6.3.7 Address resolution

9.6.3.7.1 Overview

The IPv6 layer will present the convergence layer with an IPv6 packet to be transferred. The convergence layer is responsible for determining which service node the packet should be delivered to, using the IPv6 addresses in the packet. The convergence layer shall then establish a connection to the destination if one does not already exist so that the packet can be transferred. Two classes of IPv6 addresses can be used and the following clause describes how these addresses are resolved into G.9904 EUI-48 addresses. It should be noted that IPv6 does not have a broadcast address. However broadcasting is possible using multicast all nodes addresses.

9.6.3.7.2 Unicast address

9.6.3.7.2.1 General

IPv6 unicast addresses shall be resolved into ITU-T G.9904 unicast EUI-48 addresses. The base node maintains a central database node of IPv6 addresses and EUI-48 addresses. Address resolution functions are performed by querying this database. The service node shall establish a connection to the address resolution service running on the base node, using the TYPE value TYPE_CL_IPv6_AR. No data should be passed in the connection establishment signalling. Using this connection, the service node can use two mechanisms as defined in the present Recommendation.

9.6.3.7.2.2 Address registration and deregistration

A service node uses the AR_REGISTERv6_S message to register an IPv6 address and the corresponding EUI-48 address. The base node will acknowledge an AR_REGISTERv6_B message. The service node may register multiple IPv6 addresses for the same EUI-48.

A service node uses the AR_UNREGISTERv6_S message to unregister an IPv6 address and the corresponding EUI-48 address. The base node will acknowledge an AR_UNREGISTERv6_B message.

When the address resolution connection between the service node and the base node is closed, the base node should remove all addresses associated with that connection.

9.6.3.7.2.3 Address look-up

A service node uses the AR_LOOKUPv6_S message to perform a look-up. The message contains the IPv6 address to be resolved. The base node will respond with an AR_LOOKUPv6_B message that contains an error code and, if there is no error, the EUI-48 associated with the IPv6 address. If the base node has multiple entries in its database node for the same IPv6 address, the possible EUI-48 returned is undefined.

It should be noted that, for the link local addresses, due to the fact that the EUI-48 can be obtained from the IPv6 address, the look-up can simply return this value by extracting it from the IPv6 address.

9.6.3.7.3 Multicast address

Multicast IPv6 addresses are mapped to connection handles (ConnHandle) by the convergence layer.

To join a multicast group, CL uses the MAC_JOIN.request primitive with the IPv6 address specified in the data field. A corresponding MAC_JOIN.confirm primitive will be generated by the MAC after completion of the join process. The MAC_JOIN.confirm primitive will contain the result (success/failure) and the corresponding ConnHandle to be used by the CL. The MAC layer will handle the transfer of data for this connection using the appropriate LCIDs. To leave the multicast group, the CL at the service node shall use the MAC_LEAVE.request{ConnHandle} primitive.

To send an IPv6 multicast packet, the CL will simply send the packet to the group using the allocated ConnHandle. The ConnHandle is maintained while there are more packets to be sent. However, after Tmcast_reg seconds of not sending an IPv6 multicast packet to the group, the node should release the ConnHandle by using the MAC_LEAVE.request primitive. The nominal value of Tmcast_reg is 10 minutes; however, other values may be used.

9.6.3.7.4 Retransmission of address resolution packets

The connection between the service node and the base node for address resolution is not reliable. The MAC ARQ is not used. The service node is responsible for making retransmissions if the base node does not respond in one second. It is not considered an error when the base node receives the same registration requests multiple times or is asked to remove a registration that does not exist. These conditions can be the result of retransmissions.

9.6.4 IPv6 packet transfer

For packets to be transferred, a connection needs to be established between the source and destination nodes. The IPv6 convergence layer will examine each IP packet to determine the destination EUI-48 address. If a connection to the destination has already been established, the packet is simply sent. To establish this, the convergence layer keeps a table for each connection it has with the information shown in Table 9-22. To use this table, it is first necessary to determine if the remote address is in the local subnetwork or if a gateway has to be used. The netmask associated with the local IP address is used to determine this. If the destination address is not in the local subnetwork, the address of the gateway is used instead of the destination address when the table is searched.

Table 9-22 – IPv6 convergence layer table entry

Parameter	Description
CL_IPv6_Con.Remote_IP	Remote IP address of this connection
CL_IPv6_Con.ConHandle	MAC connection handle for the connection
CL_IPv6_Con.LastUsed	Timestamp of last packet received/transmitted
CL_IPv6_Con.HC	Header compression scheme being used

The convergence layer may close a connection when it has not been used for an implementation-defined time period. When the connection is closed, the entry for the connection is removed at both ends of the connection.

When a connection to the destination does not exist, more work is necessary. The address resolution service is used to determine the EUI-48 address of the remote IP address if it is local or the gateway associated with the local address if the destination address is in another subnetwork. When the base node replies with the EUI-48 address of the destination service node, a MAC connection is established to the remote device. The TYPE value of this connection is TYPE_CL_IPv6_UNICAST. The data passed in the request message is defined in clause 9.6.8.3. The local IP address is provided so that the remote device can add the new connection to its cache of connections for sending data in the opposite direction. The use of header compression is also negotiated as part of the connection establishment. Once the connection has been established, the IP packet can be sent.

9.6.5 Segmentation and reassembly

The IPv6 convergence layer should support IPv6 packets with an MTU of 1500 bytes. This requires the use of the common part convergence sublayer segmentation and reassembly service.

9.6.6 Compression

It is assumed that any ITU-T G.9904 device is capable of LOWPAN_IPHC IPv6 header compression/decompression. It may also be capable of performing UDP compression/decompression. Thus UDP/IPv6 compression is negotiated.

No negotiation can take place for multicast packet. Nodes can only make use of mandatory compression capabilities.

Depending of the type of IPv6 address carried by the packet and the capabilities which are negotiated between the nodes involved in the data exchanges, IPv6 header compression is performed.

All the service nodes and the base node shall support IPv6 header compression using source and destination addresses stateless compression as defined in [SP 800-57]. Source and destination IPv6 addresses using stateful compression and IPv6 next header compression are negotiable.

9.6.7 Quality of service mapping

The ITU-T G.9904 MAC specifies that the contention-based access mechanism supports four priority levels (1-4). Level 1 is used for MAC signalling messages, but not exclusively so.

IPv6 packets include a traffic class field in the header to indicate the QoS the packet would like to receive. This traffic class can be used in the same way as IPv4 TOS (see [IETF RFC 791]). That is, three bits of the TOS indicate the IP precedence. The following table specifies how the IP precedence is mapped into the ITU-T G.9904 MAC priority.

Table 9-23 – Mapping IPv6 precedence to ITU-T G.9904 MAC priority

IP Precedence	MAC Priority
000 – Routine	4
001 – Priority	4
010 – Immediate	3
011 – Flash	3
100 – Flash override	2
101 – Critical	2
110 – Internetwork control	1
111 – Network control	1

NOTE – At the MAC layer level the priority as stated in the packet header field is the value assigned in this table minus 1, as the range of PKT.PRIO field is from 0 to 3.

9.6.8 Packet formats and connection data

9.6.8.1 Overview

This clause defines the format of convergence layer PDUs.

9.6.8.2 Address resolution PDU

9.6.8.2.1 General

The following PDUs are transferred over the address resolution connection between the service node and the base node. The following clauses define a number of AR.MSG values. All other values are reserved for later versions of this standard.

9.6.8.2.2 AR_REGISTERv6_S

Table 9-24 shows the address resolution register message sent from the service node to the base node.

Table 9-24 – AR_REGISTERv6_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type • For AR_REGISTERv6_S = 16
AR.IPv6	128-bits	IPv6 address to be registered
AR.EUI-48	48-bits	EUI-48 to be registered

9.6.8.2.3 AR_REGISTERv6_B

Table 9-25 shows the address resolution register acknowledgment message sent from the base node to the service node.

Table 9-25 – AR_REGISTERv6_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type • For AR_REGISTERv6_B = 17
AR.IPv6	128-bits	IPv6 address registered
AR.EUI-48	48-bits	EUI-48 registered

The AR.IPv6 and AR.EUI-48 fields are included in the AR_REGISTERv6_B message so that the service node can perform multiple overlapping registrations.

9.6.8.2.4 AR_UNREGISTERv6_S

Table 9-26 shows the address resolution unregister message sent from the service node to the base node.

Table 9-26 – AR_UNREGISTERv6_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type • For AR_UNREGISTERv6_S = 18
AR.IPv6	128-bits	IPv6 address to be unregistered
AR.EUI-48	48-bits	EUI-48 to be unregistered

9.6.8.2.5 AR_UNREGISTERv6_B

Table 9-27 shows the address resolution unregister acknowledgment message sent from the base node to the service node.

Table 9-27 – AR_UNREGISTERv6_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type • For AR_UNREGISTERv6_B = 19
AR.IPv6	128-bits	IPv6 address unregistered
AR.EUI-48	48-bits	EUI-48 unregistered

The AR.IPv6 and AR.EUI-48 fields are included in the AR_UNREGISTERv6_B message so that the service node can perform multiple overlapping unregistrations.

9.6.8.2.6 AR_LOOKUPv6_S

Table 9-28 shows the address resolution look-up message sent from the service node to the base node.

Table 9-28 – AR_LOOKUPv6_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type • For AR_LOOKUPv6_S = 20
AR.IPv6	128-bits	IPv6 address to look up

9.6.8.2.7 AR_LOOKUPv6_B

Table 9-29 shows the address resolution look-up response message sent from the base node to the service node.

Table 9-29 – AR_LOOKUPv6_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none">• For AR_LOOKUPv6_B = 21
AR.IPv6	128-bits	IPv6 address looked up
AR.EUI-48	48-bits	EUI-48 for IPv6 address
AR.Status	8-bits	Look-up status, indicating if the address was found or an error occurred. 0 = found, AR.EUI-48 valid 1 = unknown, AR.EUI-48 undefined

The look-up may fail if the requested address has not been registered. In this case, AR.Status will have a value equal to 1, and the contents of AR.EUI-48 will be undefined. The look-up is only successful when AR.Status is zero. In that case, the EUI-48 field contains the resolved address.

9.6.8.2.8 AR_MCAST_REGv6_S

Table 9-30 shows the multicast address resolution register message sent from the service node to the base node.

Table 9-30 – AR_MCAST_REGv6_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none">• For AR_MCAST_REGv6_S = 24
AR.IPv6	128-bits	IPv6 multicast address to be registered

9.6.8.2.9 AR_MCAST_REGv6_B

Table 9-31 shows the multicast address resolution register acknowledgment message sent from the base node to the service node.

Table 9-31 – AR_MCAST_REGv6_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none">• For AR_MCAST_REGv6_B = 25
AR.IPv6	128-bits	IPv6 multicast address registered
<i>Reserved</i>	2-bits	Reserved by ITU-T. Should be encoded as 0
AR.LCID	6-bits	LCID assigned to this IPv6 multicast address

The AR.IPv6 field is included in the AR_MCAST_REGv6_B message so that the service node can perform multiple overlapping registrations.

9.6.8.2.10 AR_MCAST_UNREGv6_S

Table 9-32 shows the multicast address resolution unregister messages sent from the service node to the base node.

Table 9-32 – AR_MCAST_UNREGv6_S message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none">• For AR_MCAST_UNREGv6_S = 26
AR.IPv6	128-bits	IPv6 multicast address to be unregistered

9.6.8.2.11 AR_MCAST_UNREGv6_B

Table 9-33 shows the multicast address resolution unregister acknowledgment message sent from the base node to the service node.

Table 9-33 – AR_MCAST_UNREGv6_B message format

Name	Length	Description
AR.MSG	8-bits	Address resolution message type <ul style="list-style-type: none">• For AR_MCAST_UNREGv6_B = 27
AR.IPv6	128-bits	IPv6 multicast address unregistered

The AR.IPv6 field is included in the AR_MCAST_UNREGv6_B message so that the service node can perform multiple overlapping unregistrations.

9.6.8.3 IPv6 packet format

9.6.8.3.1 General

The following PDU formats are used for transferring IPv6 packets between service nodes.

9.6.8.3.2 No negotiated header compression

When no header compression takes place, the IP packet is simply sent as it is, without any header.

Table 9-34 – IPv6 packet format without negotiated header compression

Name	Length	Description
IPv6.PKT	n-octets	The IPv6 packet

9.6.8.3.3 Header compression

When LOWPAN_IPHC1 header compression takes place, and the next header compression is negotiated, the UDP/IPv6 packet is sent as shown in Table 9-35.

Table 9-35 – UDP/IPv6 packet format with LOWPAN_IPHC1 header compression and LOWPAN_NHC

Name	Length	Description
IPv6.IPHC	2-octet	Dispatch + LOWPAN_IPHC encoding. With bit 5=1 indicating that the next is compressed, using LOWPAN_NHC format
IPv6.ncIPv6	n.m-octets	Non-compressed IPv6 fields (or elided)
IPv6.HC_UDP	1-octet	Next header encoding
IPv6.ncUDP	n.m-octets	Non-compressed UDP fields
<i>Padding</i>	0.m-octets	Padding to byte boundary
<i>IPv6.DATA</i>	n-octets	UDP data

Note that these fields are not necessarily aligned to byte boundaries. For example the IPv6.ncIPv6 field can be any number of bits. The IPv6.IPHC_UDP field follows directly afterwards, without any padding. Padding is only applied at the end of the complete compressed UDP/IPv6 header so that the UDP data is byte aligned.

When the IPv6 packet contains data other than UDP the following packet format is used as shown in Table 9-36.

Table 9-36 – IPv6 packet format with LOWPAN_IPHC negotiated header compression

Name	Length	Description
IPv6.IPHC	2-octet	HC encoding. Bits 5 contain 0 indicating the next header byte is not compressed.
IPv6.ncIPv6	n.m-octets	Non-compressed IPv6 fields
Padding	0.m-octets	Padding to byte boundary
IPv6.DATA	n-octets	IP data

9.6.8.4 Connection data

9.6.8.4.1 Overview

When a connection is established between service nodes for the transfer of IP packets, data is also transferred in the connection request packets. This data allows the negotiation of compression and notification of the IP address.

9.6.8.4.2 Connection data from the initiator

Table 9-37 shows the connection data sent by the initiator.

Table 9-37 – IPv6 connection signalling data sent by the initiator

Name	Length	Description
Reserved	6-bits	Reserved by ITU-T. Should be encoded as zero in this version of the convergence layer protocol
Data.HCNH	2-bit	Header compression negotiated <ul style="list-style-type: none"> Data.HC = 0 – No compression requested Data.HC = 1 – LOWPAN_NH Data.HC = 2 – stateful address compression Data.HC = 3 – LOWPAN_NH and stateful address compression
Data.IPv6	128-bits	IPv6 address of the initiator

If the device accepts the connection, it should copy the Data.IPv6 address into a new table entry along with the negotiated Data.HC value.

9.6.8.4.3 Connection data from the responder

Table 9-38 shows the connection data sent in response to the connection request.

Table 9-38 – IPv6 connection signalling data sent by the responder

Name	Length	Description
<i>Reserved</i>	6-bits	Reserved by ITU-T. Should be encoded as zero in this version of the convergence layer protocol
Data.HC	2-bit	Header compression negotiated <ul style="list-style-type: none"> Data.HC = 0 – No compression requested: NOTE – When stateless address compression is used, all nodes shall support it. When the stateless address compression is not used then the node notifies by this value, its compression capability. Data.HC = 1 – LOWPAN_NH Data.HC = 2 – stateful address compression Data.HC = 3 – LOWPAN_NH and stateful address compression

All nodes support stateless address compression.

The next header compression scheme and stateful address compression can only be used when it is supported by both service nodes. The responder may only set Data.HC to the same value as that received from the initiator or a value lower than the one received. When the same value is used, it indicates that the requested compression scheme has been negotiated and will be used for the connection. Setting Data.HC to a lower value allows the responder to deny the request for that header compression scheme.

9.6.9 Service access point

9.6.9.1 Overview

This clause defines the service access point used by the IPv6 layer to communicate with the IPv6 convergence layer.

9.6.9.2 Opening and closing the convergence layer

The following primitives are used to open and close the convergence layer. The convergence layer may be opened once only. The IPv6 layer may close the convergence layer when the IPv6 interface is brought down. The convergence layer will also close the convergence layer when the underlying MAC connection to the base node has been lost.

9.6.9.2.1 CL_IPv6_Establish.request

The CL_IPv6_ESTABLISH.request primitive is passed from the IPv6 layer to the IPv6 convergence layer. It is used when the IPv6 layer brings the interface up.

The semantics of this primitive are as follows:

CL_IPv6_ESTABLISH.request{}

On receiving this primitive, the convergence layer will form the address resolution connection to the base node.

9.6.9.2.2 CL_IPv6_Establish.confirm

The CL_IPv6_ESTABLISH.confirm primitive is passed from the IPv6 convergence layer to the IPv6 layer. It is used to indicate that the convergence layer is ready to access IPv6 packets to be sent to peers.

The semantics of this primitive are as follows:

CL_IPv6_ESTABLISH.confirm{}

Once the convergence layer has established all the necessary connections and is ready to transmit and receive IPv6 packets, this primitive is passed to the IPv6 layer. If the convergence layer encounters an error while opening, it responds with a CL_IPv6_RELEASE.confirm primitive, rather than a CL_IPv6_ESTABLISH.confirm.

9.6.9.2.3 CL_IPv6_Release.request

The CL_IPv6_RELEASE.request primitive is used by the IPv6 layer when the interface is put down. The convergence layer closes all connections so that no more IPv6 packets are received and all resources are released.

The semantics of this primitive are as follows:

CL_IPv6_RELEASE.request{}

Once the convergence layer has released all its connections and resources it returns a CL_IPv6_RELEASE.confirm.

9.6.9.2.4 CL_IPv6_Release.confirm

The CL_IPv6_RELEASE.confirm primitive is used by the IPv6 convergence layer to indicate to the IPv6 layer that the convergence layer has been closed. This can be as a result of a CL_IPv6_RELEASE.request primitive, a CL_IPv6_ESTABLISH.request primitive, or because the MAC layer indicates the address resolution connection has been lost, or the service node itself is no longer registered.

The semantics of this primitive are as follows:

CL_IPv6_RELEASE.confirm{Result}

The result parameter has the meanings defined in Table C.3.

9.6.9.3 Unicast address management

9.6.9.3.1 General

The primitives defined here are used for address management, i.e., the registration and unregistration of IPv6 addresses associated with this convergence layer.

When there are no IPv6 addresses associated with the convergence layer, the convergence layer will only send and receive multicast packets; unicast packets may not be sent. However, this is sufficient for various address discovery protocols to be used to gain an IPv6 address. Once an IPv6 address has been registered, the IPv6 layer can transmit unicast packets that have a source address equal to one of its registered addresses.

9.6.9.3.2 CL_IPv6_Register.request

This primitive is passed from the IPv6 layer to the IPv6 convergence layer to register an IPv6 address.

The semantics of this primitive are as follows:

CL_IPv6_REGISTER.request{ipv6, netmask, gateway}

The IPv6 address is the address to be registered.

The netmask is the network mask, used to mask the network number from the address. The netmask is used by the convergence layer to determine whether the packet should deliver directly or the gateway should be used.

The IPv6 address of the gateway, to which packets with destination addresses that are not in the same subnetwork as the local address, are to be sent.

Once the IPv6 address has been registered to the base node, a CL_IPv6_REGISTER.confirm primitive is used. If the registration fails, the CL_IPv6_RELEASE.confirm primitive will be used.

9.6.9.3.3 CL_IPv6_Register.confirm

This primitive is passed from the IPv6 convergence layer to the IPv6 layer to indicate that a registration has been successful.

The semantics of this primitive are as follows:

CL_IPv6_REGISTER.confirm{ipv6}

The IPv6 address is the address that was registered.

Once registration has been completed, the IPv6 layer may send IPv6 packets using this source address.

9.6.9.3.4 CL_IPv6_Unregister.request

This primitive is passed from the IPv6 layer to the IPv6 convergence layer to unregister an IPv6 address.

The semantics of this primitive are as follows:

CL_IPv6_UNREGISTER.request{ipv6}

The IPv6 address is the address to be unregistered.

Once the IPv6 address has been unregistered to the base node, a CL_IPv6_UNREGISTER.confirm primitive is used. If the registration fails, the CL_IPv6_RELEASE.confirm primitive will be used.

9.6.9.3.5 Unregister.confirm

This primitive is passed from the IPv6 convergence layer to the IPv6 layer to indicate that unregistration has been successful.

The semantics of this primitive are as follows:

CL_IPv6_UNREGISTER.confirm{ipv6}

The IPv6 address is the address that was unregistered.

Once unregistration has been completed, the IPv6 layer may not send IPv6 packets using this source address.

9.6.9.4 Multicast group management

9.6.9.4.1 General

This clause describes the primitives used to manage multicast groups.

9.6.9.4.2 CL_IPv6_MUL_Join.request

This primitive is passed from the IPv6 layer to the IPv6 convergence layer. It contains an IPv6 multicast address that is to be joined.

The semantics of this primitive are as follows:

CL_IPv6_MUL_JOIN.request{IPv6 }

The IPv6 address is the IPv6 multicast group that is to be joined.

When the convergence layer receives this primitive, it will arrange for IP packets sent to this group to be multicast in the ITU-T G.9904 network and receive packets using this address to be passed to the IPv6 stack. If the convergence layer cannot join the group, it uses the *CL_IPv6_MUL_LEAVE.confirm* primitive. Otherwise the *CL_IPv6_MUL_JOIN.confirm* primitive is used to indicate success.

9.6.9.4.3 CL_IPv6_MUL_Join.confirm

This primitive is passed from the IPv6 convergence layer to the IPv6. It contains a result status and an IPv6 multicast address that was joined.

The semantics of this primitive are as follows:

CL_IPv6_MUL_JOIN.confirm{IPv6}

The IPv6 address is the IPv6 multicast group that was joined. The convergence layer will start forwarding IPv6 multicast packets for the given multicast group.

9.6.9.4.4 CL_IPv6_MUL_Leave.request

This primitive is passed from the IPv6 layer to the IPv6 convergence layer. It contains an IPv6 multicast address to be left.

The semantics of this primitive are as follows:

CL_IPv6_MUL_LEAVE.request{IPv6}

The IPv6 address is the IPv6 multicast group to be left. The convergence layer will stop forwarding IPv6 multicast packets for this group and may leave the ITU-T G.9904 MAC multicast group.

9.6.9.4.5 CL_IPv6_MUL_Leave.confirm

This primitive is passed from the IPv6 convergence layer to the IPv6. It contains a result status and an IPv6 multicast address that was left.

The semantics of this primitive are as follows:

CL_IPv6_MUL_LEAVE.confirm{IPv6, Result}

The IPv6 address is the IPv6 multicast group that was left. The convergence layer will stop forwarding IPv6 multicast packets for the given multicast group.

The result takes a value from Table C.3.

This primitive can be used by the convergence layer as a result of a `CL_IPv6_MUL_JOIN.request`, `CL_IPv6_MUL_LEAVE.request` or because of an error condition resulting in the loss of the ITU-T G.9904 MAC multicast connection.

9.6.9.5 Data transfer

9.6.9.5.1 General

The following primitives are used to send and receive IPv6 packets.

9.6.9.5.2 CL_IPv6_DATA.request

This primitive is passed from the IPv6 layer to the IPv6 convergence layer. It contains one IPv6 packet to be sent.

The semantics of this primitive are as follows:

CL_IPv6_DATA.request{IPv6_PDU}

The IPv6_PDU is the IPv6 packet to be sent.

9.6.9.5.3 CL_IPv6_DATA.confirm

This primitive is passed from the IPv6 convergence layer to the IPv6 layer. It contains a status indication and an IPv6 packet that has just been sent.

The semantics of this primitive are as follows:

CL_IPv6_DATA.confirm{IPv6_PDU, Result}

The IPv6_PDU is the IPv6 packet that was to be sent.

The result value indicates whether the packet was sent or an error occurred. It takes a value from Table C.3.

9.6.9.5.4 CL_IPv6_DATA.indicate

This primitive is passed from the IPv6 convergence layer to the IPv6 layer. It contains an IPv6 packet that has just been received.

The semantics of this primitive are as follows:

CL_IPv6_DATA.indicate{IPv6_PDU}

The IPv6_PDU is the IPv6 packet that was received.

10 Management plane

10.1 Introduction

This clause specifies the management plane functionality. Figure 10-1 below highlights the position of the management plane in the overall protocol architecture.

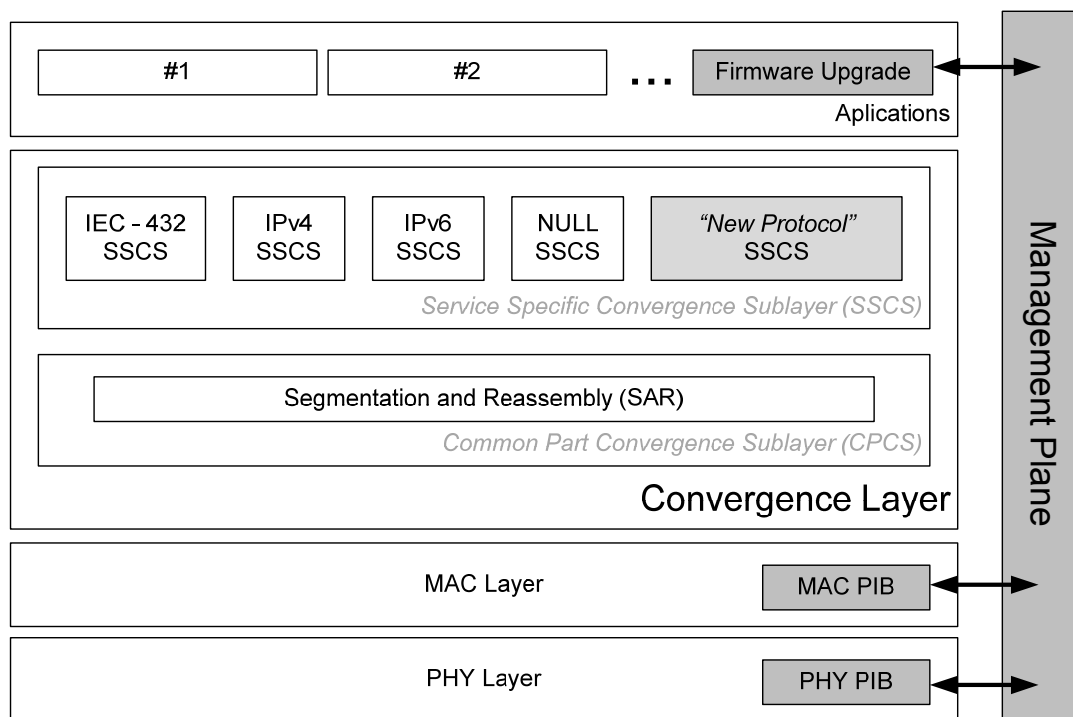


Figure 10-1 – Management plane: Introduction

All nodes shall implement the management plane functionality enumerated in this clause. The management plane enables a local or remote control entity to perform actions on a node.

The present version of this Recommendation enumerates management plane functions for node management and firmware upgrade. Future versions may include additional management functions.

- To enable access to management functions on a service node, the base node shall open a management connection after successful completion of registration (refer to clause 10.4).
- The base node may open such a connection either immediately on successful registration or some time later.
- Unicast management connections shall be identified with `CON.TYPE = TYPE_CL_MGMT`.
- Multicast management connections can also exist. At the time of writing, multicast management connections shall only be used for firmware upgrade.
- There shall be no broadcast management connection.
- In case a service node supports ARQ connections, the base node shall preferentially try to open an ARQ connection for management functions.
- Management plane functions shall use NULL SCS as specified in clause 9.3.

10.2 Node management

10.2.1 General

Node management is accomplished through a set of attributes. Attributes are defined for both PHY and MAC layers. The set of these management attributes is called a PLC information base (PIB). Some attributes are read-only while others are read-write.

PIB attribute identifiers are 16 bit values. This allows up to 65535 PIB attributes to be specified.

- PIB attribute identifier values from 0 to 32767 are open to be standardized. No proprietary attributes may have identifiers in this range.
- Values in the range 32768 to 65535 are open for vendor-specific usage.

PIB attribute identifiers in the standard range (0 to 32767) that are not specified in this version are reserved for future use.

NOTE – The PIB attribute tables below indicate the types of each attribute. For integer types the size of the integer has been specified in bits. An implementation may use a larger integer for an attribute; however, it must not use a smaller size.

10.2.2 PHY PIB attributes

10.2.2.1 General

The PHY layer implementation in each device may optionally maintain a set of attributes which provide detailed information about its working. The PHY layer attributes are part of the PLC information base (PIB).

10.2.2.2 Statistical attributes

The PHY may provide statistical information for management purposes. The next table lists the statistics that the PHY should make available to management entities across the PLME_GET primitive. TheID field in this table is the service parameter of the PLME_GET primitive specified in clause 7.10.4.9.

Table 10-1 – PHY read-only variables that provide statistical information

Attribute name	Size (in bits)	Id	Description
phyStatsCRCIncorrectCount	16	0x00A0	Number of bursts received on the PHY layer for which the CRC was incorrect.
phyStatsCRCFailCount	16	0x00A1	Number of bursts received on the PHY layer for which the CRC was correct, but the <i>Protocol</i> field of the PHY header had an invalid value. This count would reflect the number of times corrupt data was received and the CRC calculation failed to detect it.
phyStatsTxDropCount	16	0x00A2	Number of times that the PHY layer received new data to transmit (PHY_DATA.request) and had to either overwrite existing data in its transmit queue or drop the data in new requests due to a full queue.
phyStatsRxDropCount	16	0x00A3	Number of times when the PHY layer received new data on the channel and had to either overwrite existing data in its receive queue or drop the newly received data due to full queue.
phyStatsRxTotalCount	32	0x00A4	Total number of PPDUs correctly decoded. Useful for PHY layer test cases, to estimate the FER.

Table 10-1 – PHY read-only variables that provide statistical information

Attribute name	Size (in bits)	Id	Description
phyStatsBlkAvgEvm	16	0x00A5	Exponential moving average of the EVM over the past 16 PPDU's, as returned by the PHY_SNR primitive. Note that the PHY_SNR primitive returns a 3-bit number in dB scale. So first each 3-bit dB number is converted to linear scale (number k goes to $2^{(k/2)}$), yielding a 7 bit number with 3 fractional bits. The result is just accumulated over 16 PPDU's and reported.
phyEmaSmoothing	8	0x00A8	Smoothing factor divider for values that are updated as the exponential moving average (EMA). Next value is $V_{next} = S \times NewSample + (1-S) \times V_{prev}$ Where $S = 1 / (2^{phyEMASmoothing})$

10.2.2.3 Implementation attributes

It is possible to implement PHY functions conforming to this Recommendation in multiple ways. The multiple implementation options provide some degree of unpredictability for the MAC layers. PHY implementations may optionally provide specific information on parameters which are of interest to the MAC across the PLME_GET primitive. A list of such parameters which maybe queried across the PLME_GET primitives by the MAC is provided in Table 10-2. All of the attributes listed in Table 10-2 are implementation constants and shall not be changed.

Table 10-2 – PHY read-only parameters, providing information on specific implementation

Attribute name	Size (in bits)	Id	Description
phyTxQueueLen	10	0x00B0	Number of concurrent MPDU's that the PHY transmit buffers can hold.
phyRxQueueLen	10	0x00B1	Number of concurrent MPDU's that the PHY receive buffers can hold.
phyTxProcessingDelay	20	0x00B2	Time elapsed from the instance when data is received on the MAC-PHY communication interface to the time when it is put on the physical channel. This shall not include communication delay over the MAC-PHY interface. Value of this attribute is in unit of microseconds.
phyRxProcessingDelay	20	0x00B3	Time elapsed from the instance when data is received on the physical channel to the time when it is made available to the MAC across the MAC-PHY communication interface. This shall not include communication delay over the MAC-PHY interface. Value of this attribute is in unit of microseconds.
phyAgcMinGain	8	0x00B4	Minimum gain for the AGC \leq 0dB.

Table 10-2 – PHY read-only parameters, providing information on specific implementation

Attribute name	Size (in bits)	Id	Description
phyAgcStepValue	3	0x00B5	Distance between steps in dB ≤ 6dB.
phyAgcStepNumber	8	0x00B6	Number of steps so that $\text{phyAgcMinGain} + (\text{phyAgcStepNumber} - 1) \times \text{phyAgcStepValue} \geq 21$ dB.

10.2.3 MAC PIB attributes

10.2.3.1 General

NOTE – Note that the "M"(Mandatory) column in the tables below specifies if the PIB attributes are mandatory for all devices (both service node and base node, specified as "All"), only for service nodes ("SN"), only for base nodes ("BN") or not mandatory at all ("No").

10.2.3.2 MAC variable attributes

MAC PIB variables include the set of PIB attributes that influence the functional behaviour of an implementation. These attributes may be defined externally to the MAC, typically by the management entity and implementations may allow changes to their values during normal running, i.e., even after the device start-up sequence has been executed.

An external management entity can have access to these attributes through the MLME_GET (clause 8.5.5.7) and MLME_SET (clause 8.5.5.9) set of primitives. TheID field in the following table would be the PIBAttribute that needs to be passed to the MLME SAP while working on these parameters.

Table 10-3 – Table of MAC read-write variables

Attribute name	Id	Type	M	Valid range	Description	Def.
macMinSwitchSearch Time	0x0010	Integer8	No	16-32 seconds	Minimum time for which a service node in the disconnected status should scan the channel for beacons before it can broadcast PNPDU. This attribute is not maintained in base nodes.	24
macMaxPromotionPdu	0x0011	Integer8	No	1-4	Maximum number of PNPDU that may be transmitted by a service node in a period of <i>macPromotionPduTxPeriod</i> seconds. This attribute is not maintained in base nodes.	2

Table 10-3 – Table of MAC read-write variables

Attribute name	Id	Type	M	Valid range	Description	Def.
macPromotionPduTx Period	0x0012	Integer8	No	2-8 seconds	Time quantum for limiting a number of PNPDU's transmitted from a service node. No more than <i>macMaxPromotionPdu</i> may be transmitted in a period of <i>macPromotionPduTxPeriod</i> seconds.	5
macBeaconsPerFrame	0x0013	Integer8	BN	1-5	Maximum number of beacon slots that may be provisioned in a frame. This attribute is maintained in base nodes.	5
macSCPMaXTxAttempts	0x0014	Integer8	No	2-5	Number of times the CSMA algorithm would attempt to transmit requested data when a previous attempt was withheld due to the PHY indicating channel busy.	5
macCtlReTxTimer	0x0015	Integer8	No	2-20 seconds	Number of seconds for which a MAC entity waits for acknowledgement of receipt of the MAC control packet from its peer entity. On expiry of this time, the MAC entity may retransmit the MAC control packet.	15
macMaxCtlReTx	0x0018	Integer8	No	3-5	Maximum number of times a MAC entity will try to retransmit an unacknowledged MAC control packet. If the retransmit count reaches this maximum, the MAC entity shall abort further attempts to transmit the MAC control packet.	3

Table 10-3 – Table of MAC read-write variables

Attribute name	Id	Type	M	Valid range	Description	Def.
macEMASmoothing	0x0019	Integer8	All	0-7	Smoothing factor divider for values that are updated as the exponential moving average (EMA). Next value is $V_{next} = S \times NewSample + (1 - S) \times V_{prev}$ Where $S = 1 / (2^{\text{macEMASmoothing}})$.	3

Table 10-4 – Table of MAC read-only variables

Attribute name	Id	Type	M	Valid range	Description	Def.
macSCPRBO	0x0016	Integer8	No	1-15 symbols	Random backoff period for which an implementation should delay the start of channel-sensing iterations when attempting to transmit data in the SCP. This is a 'read-only' attribute.	–
macSCPChSenseCount	0x0017	Integer8	No	2-5	Number of times for which an implementation has to perform channel-sensing. This is a 'read-only' attribute.	–

10.2.3.3 Functional attributes

Some PIB attributes belong to the functional behaviour of the MAC. They provide information on specific aspects. A management entity can only read their present value using the MLME_GET primitives. The value of these attributes cannot be changed by a management entity through the MLME_SET primitives.

TheID field in the table below would be the *PIBAttribute* that needs to be passed MLME_GET SAP for accessing the value of these attributes.

Table 10-5 – Table of MAC read-only variables that provide functional information

Attribute name	Id	Type	M	Valid range	Description
macLNID	0x0020	Integer16	SN	0-16383	LNID allocated to this node at time of its registration.
MacLSID	0x0021	Integer8	SN	0-255	LSID allocated to this node at time of its promotion. This attribute is not maintained if a node is in a terminal functional state.
MacSID	0x0022	Integer8	SN	0-255	SID of the switch node through which this node is connected to the subnetwork. This attribute is not maintained in a base node.
MacSNA	0x0023	EUI-48	SN		Subnetwork address to which this node is registered. The base node returns the SNA it is using.
MacState	0x0024	Enumerate	SN		Present functional state of the node
				0	DISCONNECTED
				1	TERMINAL
				2	SWITCH
				3	BASE
MacSCPLength	0x0025	Integer16	SN		The SCP length, in symbols, in the present frame.
MacNodeHierarchy Level	0x0026	Integer8	SN	0-63	Level of this node in subnetwork hierarchy.
MacBeaconSlotCount	0x0027	Integer8	SN	0-7	Number of beacon slots provisioned in the present frame structure.
macBeaconRxSlot	0x0028	Integer8	SN	0-7	Beacon slot on which this device's switch node transmits its beacon. This attribute is not maintained in a base node.
MacBeaconTxSlot	0x0029	Integer8	SN	0-7	Beacon slot in which this device transmits its beacon. This attribute is not maintained in service nodes that are in a terminal functional state.
MacBeaconRxFrequency	0x002A	Integer8	SN	0-31	Number of frames between the reception of two successive beacons. A value of 0x0 indicates beacons are received in every frame. This attribute is not maintained in base nodes.

Table 10-5 – Table of MAC read-only variables that provide functional information

Attribute name	Id	Type	M	Valid range	Description
MacBeaconTxFrequency	0x002B	Integer8	SN	0-31	Number of frames between transmissions of two successive beacons. A value of 0x0 indicates beacons are transmitted in every frame. This attribute is not maintained in service nodes that are in a terminal functional state.
MacCapabilities	0x002C	Integer16	All	Bit-map	<p>Bitmap of MAC capabilities of a given device. This attribute shall be maintained on all devices. Bits in sequence of right-to-left shall have the following meaning:</p> <ul style="list-style-type: none"> Bit 0: switch capable; Bit 1: packet aggregation; Bit 2: contention-free period; Bit 3: direct connection; Bit 4: multicast; Bit 5: PHY robustness management; Bit 6: ARQ; Bit 7: reserved by ITU-T; Bit 8: direct connection switching; Bit 9: multicast switching capability; Bit 10: PHY robustness management switching capability; Bit 11: ARQ buffering switching capability; Bits 12 to 15: reserved by ITU-T.

10.2.3.4 Statistical attributes

The MAC layer shall provide statistical information for management purposes. Table 10-6 lists the statistics that the MAC shall make available to management entities across the MLME_GET primitive.

TheID field in the table below would be the *PIBAttribute* that needs to be passed to the MLME_GET SAP for accessing the value of these attributes.

Table 10-6 – Table of MAC read-only variables that provide statistical information

Attribute name	Id	M	Type	Description
macTxDataPktCount	0x0040	No	Integer32	Count of successfully transmitted MSDUs.
MacRxDataPktCount	0x0041	No	Integer32	Count of successfully received MSDUs whose destination address was this node.
MacTxCtrlPktCount	0x0042	No	Integer32	Count of successfully transmitted MAC control packets.
MacRxCtrlPktCount	0x0043	No	Integer32	Count of successfully received MAC control packets whose destination address was this node.
MacCSMAFailCount	0x0044	No	Integer32	Count of failed CSMA transmitted attempts.
MacCSMAChBusyCount	0x0045	No	Integer32	Count of number of times this node had to back off SCP transmission due to a channel busy state.

10.2.3.5 MAC list attributes

The MAC layer shall make certain lists available to the management entity across the MLME_LIST_GET primitive. These lists are given in Table 10-7. Although a management entity can read each of these lists, it cannot change the contents of any of them.

TheID field in the table below would be the *PIBListAttribute* that needs to be passed to the MLME_LIST_GET primitive for accessing the value of these attributes.

Table 10-7 – Table of read-only lists made available by MAC layer through the management interface

List attribute Name	Id	M	Description		
macListRegDevices	0x0050	BN	List of registered devices. This list is maintained by the base node only. Each entry in this list shall comprise the following information.		
			Entry element	Type	Description
			regEntryID	EUI-48	EUI-48 of the registered node
			regEntryLNID	Integer16	LNID allocated to this node
			regEntryState	TERMINAL=1, SWITCH=2	Functional state of this node
			regEntryLSID	Integer16	SID allocated to this node
			regEntrySID	Integer16	SID of switch through which this node is connected
			regEntryLevel	Integer8	Hierarchy level of this node

Table 10-7 – Table of read-only lists made available by MAC layer through the management interface

List attribute Name	Id	M	Description		
			regEntryTCap	Integer8	Bitmap of MAC capabilities of terminal functions in this device Bits in sequence of right-to-left shall have the following meaning: Bit 0: switch capable; Bit 1: packet aggregation; Bit 2: contention-free period; Bit 3: direct connection; Bit 4: multicast; Bit 5: PHY robustness management; Bit 6: ARQ; Bit 7: reserved by ITU-T.
			regEntrySwCap	Integer8	Bitmap of MAC switching capabilities of this device Bits in sequence of right-to-left shall have the following meaning: Bit 0: direct connection switching capability; Bit 1: multicast switching; Bit 2:PHY robustness management switching capability; Bit 3:ARQ buffering switching capability; Bit 4 to 7: reserved by ITU-T.
macListActiveConn	0x0051	BN	List of active non-direct connections. This list is maintained by the base node only.		
			Entry element	Type	Description
			connEntrySID	Integer16	SID of switch through which the service node is connected.
			connEntryLNID	Integer16	NID allocated to service node
			connEntryLCID	Integer8	LCID allocated to this connection
			connEntryID	EUI-48	EUI-48 of service node

Table 10-7 – Table of read-only lists made available by MAC layer through the management interface

List attribute Name	Id	M	Description		
macListMcast Entries	0x0052	No	List of entries in multicast switching table. This list is not maintained by service nodes in a terminal functional state.		
			Entry element	Type	Description
			mcastEntryLCID	Integer8	LCID of the multicast group
			mcastEntry Members	Integer16	Number of child nodes (including the node itself) that are members of this group.
macListSwitchTable	0x0053	SN	List the switch table. This list is not maintained by service nodes in a terminal functional state.		
			Entry element	Type	Description
			stblEntryLSID	Integer16	SID of attached switch node
macListDirectConn	0x0054	No	List of direct connections that are active. This list is maintained only in the base node.		
			Entry element	Type	Description
			dconnEntrySrcSID	Integer16	SID of switch through which the source service node is connected
			dconEntrySrcLNID	Integer16	NID allocated to the source service node
			dconnEntrySrcLCID	Integer8	LCID allocated to this connection at the source
			dconnEntrySrcID	EUI-48	EUI-48 of source service node
			dconnEntryDstSID	Integer16	SID of switch through which the destination service node is connected
			dconnEntryDstLNID	Integer16	NID allocated to the destination service node
			dconnEntryDstLCID	Integer8	LCID allocated to this connection at the destination
			dconnEntryDstID	EUI-48	EUI-48 of destination service node
			dconnEntryDSID	Integer16	SID of switch that is the direct switch
			dconnEntryDID	EUI-48	EUI-48 of direct switch

Table 10-7 – Table of read-only lists made available by MAC layer through the management interface

List attribute Name	Id	M	Description		
macListDirectTable	0x0055	No	List the direct switch table		
			Entry element	Type	Description
			dconnEntrySrcSID	Integer16	SID of switch through which the source service node is connected
			dconEntrySrcLNID	Integer16	NID allocated to the source service node
			dconnEntrySrcLCID	Integer8	LCID allocated to this connection at the source
			dconnEntryDstSID	Integer16	SID of switch through which the destination service node is connected
			dconnEntryDstLNID	Integer16	NID allocated to the destination service node
			dconnEntryDstLCID	Integer8	LCID allocated to this connection at the destination
			dconnEntryDID	EUI-48	EUI-48 of direct switch
macListAvailable Switches	0x0056	SN	List of switch nodes whose beacons are received.		
			Entry Element	Type	Description
			slistEntrySNA	EUI-48	EUI-48 of the subnetwork
			slistEntryLSID	Integer16	SID of this switch
			slistEntryLevel	Integer8	Level of this switch in subnetwork hierarchy
			slistEntryRxLvl	Integer8EMA	Received signal level for this switch
			slistEntryRxSNR	Integer8 EMA	Signal to noise ratio for this switch
macListPhyComm	0x0057	All	List of PHY communication parameters. This table is maintained in every node. For terminal nodes it contains only one entry for the switch the node is connected through. For other nodes this contains also entries for every directly connected child node.		
			Entry element	Type	Description
			phyCommEUI	EUI-48	EUI-48 of the other device
			phyCommTxPwr	Integer8	Tx power of GPDU packets sent to the device
			phyCommTxCod	Integer8	Tx coding of GPDU packets sent to the device
			phyCommRxCod	Integer8	Rx coding of GPDU packets received from the device

Table 10-7 – Table of read-only lists made available by MAC layer through the management interface

List attribute Name	Id	M	Description		
			phyCommRxLvl	Integer8EMA	Rx power level of GPDU packets received from the device
			phyCommSNR	Integer8EMA	SNR of GPDU packets received from the device
			phyCommTxPwr Mod	Integer8	Number of times the Tx power was modified
			phyCommTxCod Mod	Integer8	Number of times the Tx coding was modified
			phyCommRxCod Mod	Integer8	Number of times the Rx coding was modified

10.2.3.6 Action PIB attributes

Some of the conformance tests require triggering certain actions on service nodes. The following table lists the set of action attributes that need to be supported by all implementations.

Table 10-8 – Action PIB attributes

Attribute name	Id	M	Size (in bits)	Description
MACActionTxData	0x0060	SN	8	Total number of PPDUs correctly decoded. Useful for the PHY layer to estimate FER.
MACActionConnClose	0x0061	SN	8	Trigger to close one of the open connections
MACActionRegReject	0x0062	SN	8	Trigger to reject incoming registration request
MACActionProReject	0x0063	SN	8	Trigger to reject incoming promotion request
MACActionUnregister	0x0064	SN	8	Trigger to unregister from the subnetwork

10.2.4 Application PIB attributes

The following PIB attributes are used for general administration and maintenance of an ITU-T G.9904 compliant device. These attributes do not affect the communication functionality, but enable easier administration.

These attributes shall be supported by both base node and service node devices.

Table 10-9 – Applications PIB attributes

Attribute Name	Size (in bits)	Id	Description
AppFwVersion	128	0x0075	Textual description of firmware version running on device
AppVendorId	16	0x0076	PRIME Alliance-assigned unique vendor identifier
AppProductId	16	0x0077	Vendor-assigned unique identifier for specific product

10.3 Firmware upgrade

10.3.1 General

The present clause specifies firmware upgrade. Devices supporting ITU-T G.9904 may have several firmware inside them, at least one supporting the application itself, and one related to the ITU-T G.9904 protocol. Although it is possible that the application can perform the firmware upgrade of all the firmware images of the device, for instance DLMS/COSEM image transfer, using COSEM image transfer object, supporting ITU-T G.9904 firmware upgrade is mandatory in order to process the ITU-T G.9904 firmware upgrade independently of the application.

10.3.2 Requirements and features

This clause specifies the firmware upgrade application, which is unique and mandatory for base nodes and service nodes.

The most important features of the firmware upgrade mechanism are listed below. See the following clauses for more information. The following refer to the FU mechanism:

- It shall be a part of the management plane and therefore use the NULL SSCS, as specified in clause 9.3.
- It is able to work in unicast (default mode) and multicast (optional mode). The control messages are always sent using unicast connections, whereas data can be transmitted using both unicast and multicast. No broadcast should be used to transmit data.
- It may change the data packet sizes according to the channel conditions. The packet size will not be changed during the download process.
- It is able to request basic information from the service nodes at any time, such as device model, firmware version and FU protocol version.
- It shall be able to be aborted at any time.
- It shall check the integrity of the downloaded FW after completing the reception. In case of failure, the firmware upgrade application shall request a new retransmission.
- The new firmware shall be executed in the service nodes only if they are commanded to do so. The FU application shall have to be able to set the moment when the reset takes place.
- Must be able to reject the new firmware after a "test" period and switch to the old version. The duration of this test period has to be fixed by the FU mechanism.

10.3.3 General description

10.3.3.1 General

The firmware upgrade mechanism is able to work in unicast and multicast modes. All control messages are sent using unicast connections, whereas the data can be sent via unicast (by default) or multicast (only if supported by the manufacturer). Note that in order to ensure correct reception of the FW when service nodes from different vendors are upgraded, data packets shall not be sent via broadcast. Only unicast and multicast are allowed. A node will reply only to messages sent via unicast. See clause 10.3.5 for a detailed description of the control and information messages used by the FU mechanism.

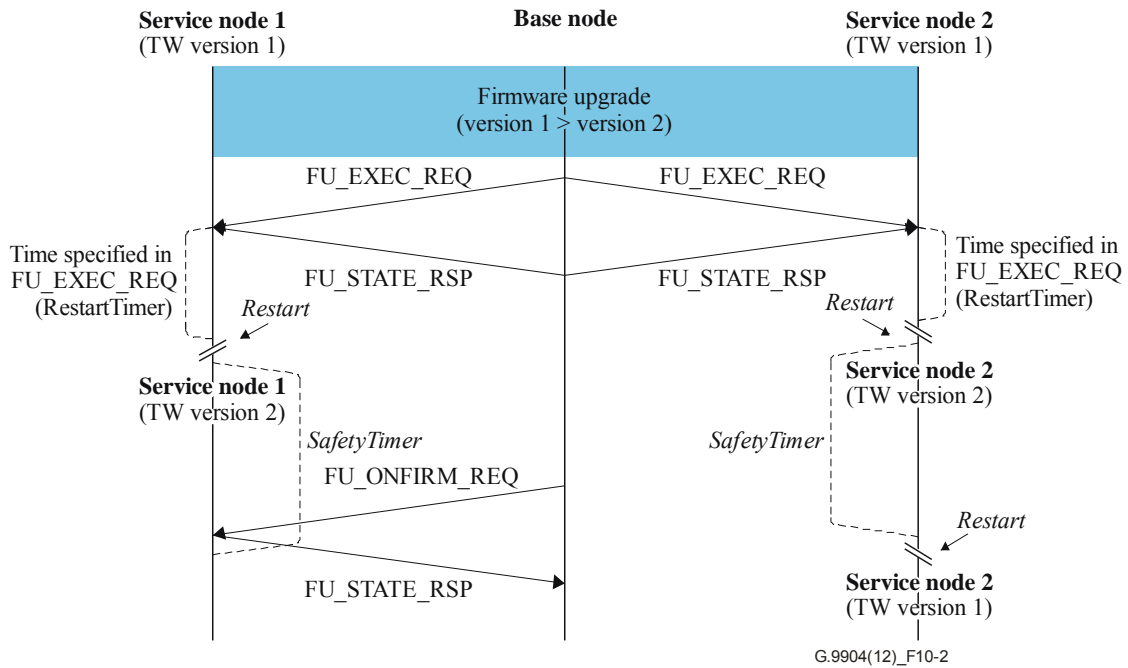
The unicast and multicast connections are set up by the base node. In the case of supporting multicast, the base node shall request the nodes from a specific vendor to join a specific multicast group, which is exclusively created to perform the firmware upgrade and is removed after finishing it.

As said before, it is up to the vendor to use unicast or multicast for transmitting the data. In the case of unicast data transmission, please note that the use of ARQ is an optional feature. Some examples showing the traffic between the base node and the service nodes in unicast and multicast are provided in clause 10.3.6.

After completing the firmware download, each service node is committed by the base node to perform an integrity check on it. The firmware download will be restarted if the firmware image results are corrupted. In another case, the service nodes will wait until they are commanded by the base node to execute the new firmware.

The FU mechanism can set up the instant when the recently downloaded firmware is executed on the service nodes. Thus, the base node can choose to restart all nodes at the same time or in several steps. After restart, each service node runs the new firmware for a time period specified by the FU mechanism. If this period expires without receiving any confirmation from the base node, or the base node decides to abort the upgrade process, the service nodes will reject the new firmware and switch to the old version. In any other case (a confirmation message is received) the service nodes will consider the new firmware as the only valid version and delete the old one.

This is done in order to leave an "open back-door" in case the new firmware is defect or corrupt. Please note that the service nodes are not allowed to discard any of the stored firmware versions until the final confirmation from the base node arrives or until the safety time period expires. The two last firmware upgrade steps explained above are shown in clause 10.3.5. See clause 10.3.5.3 for a detailed description of the control messages.



NOTE – In normal circumstances, both service nodes should either accept or reject the new firmware version. Both possibilities are shown above simultaneously for academic purposes.

Figure 10-2 – Restarting the nodes and running the new firmware

10.3.3.2 Segmentation

The firmware image is the information to be transferred, in order to process a firmware upgrade. The size of the firmware image will be called "*ImageSize*", and is measured in bytes. This image is divided into smaller elements called pages that are easier to be transferred in packets. The "*PageSize*" may be one of the following: 32 bytes, 64 bytes, 128 bytes or 192 bytes. This implies that the number of pages in a firmware image is calculated by the following formula:

$$PageCount = \left\lceil \frac{ImageSize}{PageSize} \right\rceil + 1$$

Every page will have a size specified by *PageSize*, except the last one that will contain the remaining bytes up to *ImageSize*.

The *PageSize* is configured by the base node and notified during the initialization of the firmware upgrade process, and imposes a condition in the size of the packets being transferred by the protocol.

10.3.4 Firmware upgrade PIB attributes

The following PIB attributes shall be supported by service nodes to support the firmware download application.

Table 10-10 – FU PIB attributes

Attribute name	Size (in bits)	Id	Description
AppFwdlRunning	16	0x0070	Indicates if a firmware download is in progress or not. 0 = No firmware download 1 = Firmware download in progress
AppFwdlRxPktCount	16	0x0071	Count of firmware download packets that have been received until the time of query

10.3.5 State machine

10.3.5.1 General

A service node using the firmware upgrade service will be in one of five possible states: *Idle*, *Receiving*, *Complete*, *Countdown* and *Upgrade*. These states, the events triggering them and the resulting actions/output messages are detailed below.

Table 10-11 – FU state machine

FU State	Description	Event	Output (or action to be performed)	Next state
<i>Idle</i>	The FU application is doing nothing.	Receive FU_INFO_REQ	FU_INFO_RSP.	<i>Idle</i>
		Receive FU_STATE_REQ	FU_STATE_RSP (.State = 0).	<i>Idle</i>
		Receive FU_MISS_REQ	FU_STATE_RSP (.State = 0).	<i>Idle</i>
		Receive FU_CRC_REQ	FU_STATE_RSP (.State = 0).	<i>Idle</i>
		Receive FU_INIT_REQ	FU_STATE_RSP (.State = 1).	<i>Receiving</i>
		Receive FU_DATA	(ignore).	<i>Idle</i>
		Receive FU_EXEC_REQ	FU_STATE_RSP (.State = 0).	<i>Idle</i>
		Receive FU_CONFIRM_REQ	FU_STATE_RSP (.State = 0).	<i>Idle</i>
		Receive FU_KILL_REQ	FU_STATE_RSP (.State = 0).	<i>Idle</i>

Table 10-11 – FU state machine

FU State	Description	Event	Output (or action to be performed)	Next state
<i>Receiving</i>	The FU application is receiving the firmware image.	Complete FW received and CRC OK	(if CRC of the complete image is OK, switch to <i>Complete</i> without sending any additional messages)	<i>Complete</i>
		Receive FU_INFO_REQ	FU_INFO_RSP.	<i>Receiving</i>
		Receive FU_STATE_REQ	FU_STATE_RSP (.State = 1).	<i>Receiving</i>
		Receive FU_MISS_REQ	FU_MISS_LIST or FU_MISS_BITMAP.	<i>Receiving</i>
		Receive FU_CRC_REQ	FU_CRC_RSP (FU_STATE_RSP if the Bitmap is not complete)	<i>Receiving</i>
		Receive FU_INIT_REQ	FU_STATE_RSP (.State = 1)	<i>Receiving</i>
		Receive FU_DATA	(receiving data, normal behaviour).	<i>Receiving</i>
		Receive FU_EXEC_REQ	FU_STATE_RSP (.State = 1).	<i>Receiving</i>
		Receive FU_CONFIRM_REQ	FU_STATE_RSP (.State = 1).	<i>Receiving</i>
		Receive FU_KILL_REQ	FU_STATE_RSP (.State = 0); (switch to <i>Idle</i>).	<i>Idle</i>

Table 10-11 – FU state machine

FU State	Description	Event	Output (or action to be performed)	Next state
Complete	Upgrade completed, image integrity OK, the service node is waiting to reboot with the new FW version.	Receive FU_INFO_REQ	FU_INFO_RSP.	<i>Complete</i>
		Receive FU_STATE_REQ	FU_STATE_RSP (.State = 2).	<i>Complete</i>
		Receive FU_MISS_REQ	FU_STATE_RSP (.State = 2).	<i>Complete</i>
		Receive FU_CRC_REQ	FU_STATE_RSP (.State = 2).	<i>Complete</i>
		Receive FU_INIT_REQ	FU_STATE_RSP (.State = 2).	<i>Complete</i>
		Receive FU_DATA	(ignore).	<i>Complete</i>
		Receive FU_EXEC_REQ with RestartTimer != 0	FU_STATE_RSP (.State = 3).	<i>Countdown</i>
		Receive FU_EXEC_REQ with RestartTimer = 0	FU_STATE_RSP (.State = 4).	<i>Upgrade</i>
		Receive FU_CONFIRM_REQ	FU_STATE_RSP (.State = 2).	<i>Complete</i>
		Receive FU_KILL_REQ	FU_STATE_RSP (.State = 0); (switch to <i>Idle</i>).	<i>Idle</i>
Countdown	Waiting until RestartTimer expires.	RestartTimer expires	(switch to <i>Upgrade</i>).	<i>Upgrade</i>
		Receive FU_INFO_REQ	FU_INFO_RSP.	<i>Countdown</i>
		Receive FU_STATE_REQ	FU_STATE_RSP (.State = 3).	<i>Countdown</i>
		Receive FU_MISS_REQ	FU_STATE_RSP (.State = 3).	<i>Countdown</i>
		Receive FU_CRC_REQ	FU_STATE_RSP (.State = 3).	<i>Countdown</i>
		Receive FU_INIT_REQ	FU_STATE_RSP (.State = 3).	<i>Countdown</i>
		Receive FU_DATA	(ignore).	<i>Countdown</i>
		Receive FU_EXEC_REQ with RestartTimer != 0	FU_STATE_RSP (.State = 3); (update RestartTimer and SafetyTimer).	<i>Countdown</i>
		Receive FU_EXEC_REQ with RestartTimer = 0	FU_STATE_RSP (.State = 4); (update RestartTimer and SafetyTimer).	<i>Upgrade</i>
		Receive FU_CONFIRM_REQ	FU_STATE_RSP (.State = 3).	<i>Countdown</i>
		Receive FU_KILL_REQ	FU_STATE_RSP (.State = 0); (switch to <i>Idle</i>).	<i>Idle</i>

Table 10-11 – FU state machine

FU State	Description	Event	Output (or action to be performed)	Next state
<i>Upgrade</i>	The FU mechanism reboots using the new FW image and tests it for <i>SafetyTimer</i> seconds.	<i>SafetyTimer</i> expires	FU_STATE_RSP (.State = 0); (switch to <i>Idle</i> , FW rejected).	<i>Idle</i>
		Receive FU_INFO_REQ	FU_INFO_RSP.	<i>Upgrade</i>
		Receive FU_STATE_REQ	FU_STATE_RSP (.State = 4).	<i>Upgrade</i>
		Receive FU_MISS_REQ	FU_STATE_RSP (.State = 4).	<i>Upgrade</i>
		Receive FU_CRC_REQ	FU_STATE_RSP (.State = 4).	<i>Upgrade</i>
		Receive FU_INIT_REQ	FU_STATE_RSP (.State = 4).	<i>Upgrade</i>
		Receive FU_DATA	(ignore).	<i>Upgrade</i>
		Receive FU_EXEC_REQ	FU_STATE_RSP (.State = 0).	<i>Upgrade</i>
		Receive FU_CONFIRM_REQ	FU_STATE_RSP (.State = 0); (switch to <i>Idle</i> , FW accepted).	<i>Idle</i>
		Receive FU_KILL_REQ	FU_STATE_RSP (.State = 0); (switch to <i>Idle</i> , FW rejected).	<i>Idle</i>

The state diagram is represented below. Please note that only the most relevant events are shown in the state transitions. See clause 10.3.5 for a detailed description of each state's behaviour and the events and actions related to them. A short description of each state is provided in clause 10.3.5.2.

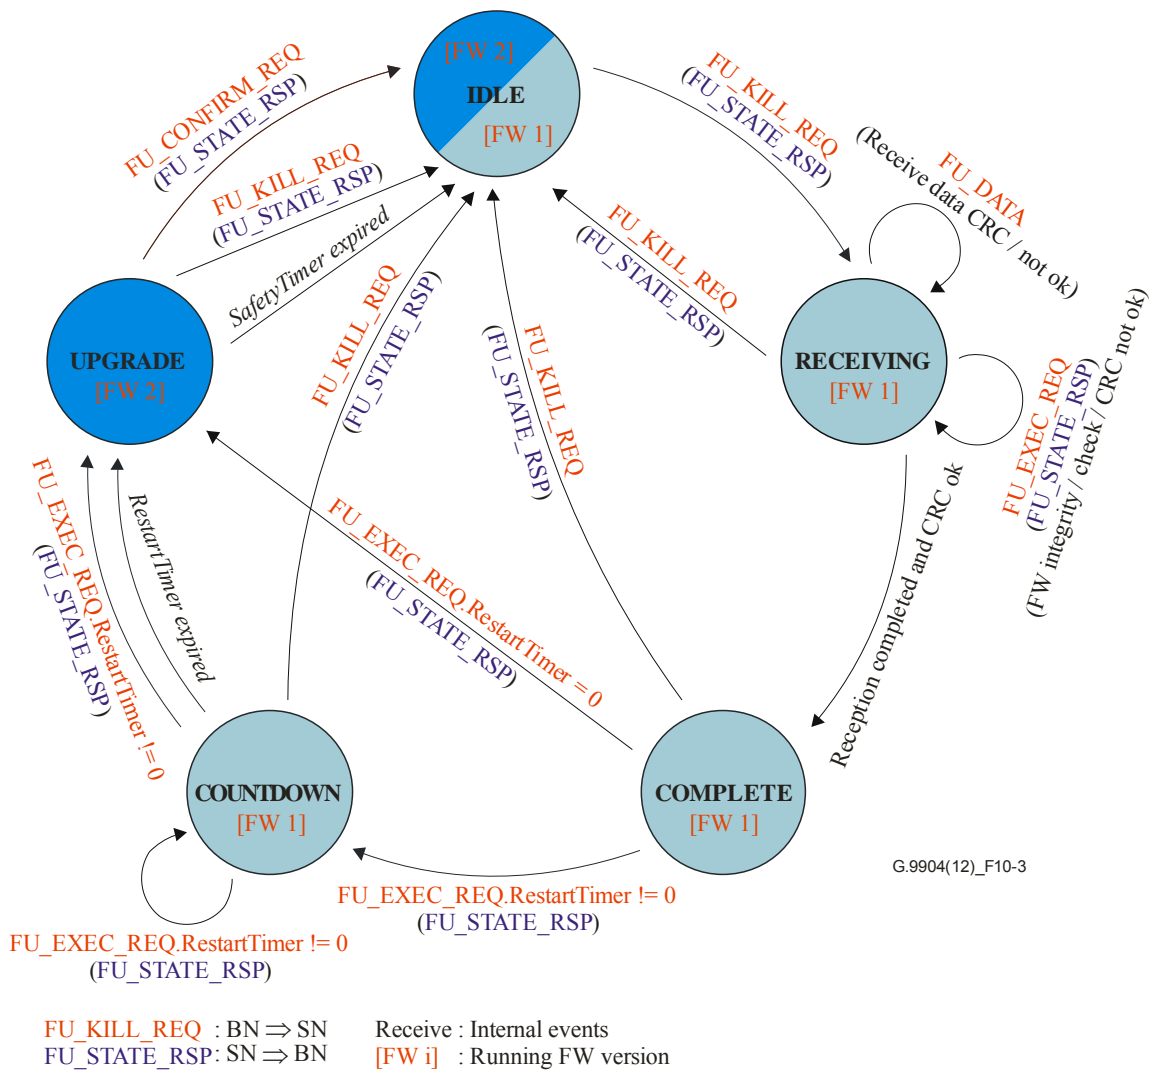


Figure 10-3 – Firmware upgrade mechanism, state diagram

10.3.5.2 State description

10.3.5.2.1 Idle

The service nodes are in "Idle" state when they are not performing a firmware upgrade. The reception of an `FU_INIT_REQ` message is the only event that forces the service node to switch to the next state ("Receiving"). `FU_KILL_REQ` aborts the upgrade process and forces the service nodes to switch from any state to "Idle".

10.3.5.2.2 Receiving

The service nodes receive the firmware image via `FU_DATA` messages. Once the download is complete, the integrity of the image is checked by the base node using `FU_CRC_REQ` and the service node responds with `FU_CRC_RSP`. This final CRC on the complete FW image is mandatory. If the CRC results are OK, the service node responds with `FU_CRC_RSP` and then switches to "Complete" state. If the CRC is wrong, the service node reports this to the base node via `FU_CRC_RSP`, drops the complete FW image, updates the bitmap accordingly and waits for packet retransmission.

Please remember that the service node will change from "Receiving" to "Complete" state only if the complete FW has been downloaded and the CRC has been successful.

10.3.5.2.3 Complete

A service node in "Complete" state waits until the reception of an FU_EXEC_REQ message. The service node may switch either to "Countdown" or "Upgrade" depending on the field *RestartTimer*, which specifies in which instant the service node has to reboot using the new firmware. If *RestartTimer* = 0, the service node immediately switches to "Upgrade"; else, the service node switches to "Countdown".

10.3.5.2.4 Countdown

A service node in "Countdown" state waits a period of time specified in the *RestartTimer* field of a previous FU_EXEC_REQ message. When this timer expires, it automatically switches to "Upgrade".

FU_EXEC_REQ can be used in "Countdown" state to reset *RestartTimer* and *SafetyTimer*. In this case, both timers have to be specified in FU_EXEC_REQ because both will be overwritten. Note that it is possible to force the node to immediately switch from the "Countdown" to "Upgrade" state setting *RestartTimer* to zero.

10.3.5.2.5 Upgrade

A service node in "Upgrade" state shall run the new firmware during a time period specified in FU_EXEC_REQ.SafetyTimer. If it does not receive any confirmation at all before this timer expires (or if it receives an FU_KILL_REQ message), the service node discards the new FW, reboots with the old version and switches to "Idle" state. In any other case it discards the old FW version and switches to "Idle" state.

10.3.5.3 Control packets

10.3.5.3.1 FU_INIT_REQ

The base node sends this packet in order to configure a service node for the Firmware upgrade. If the service node is in "Idle" state, it will change its state from "Idle" to "Receiving" and will answer with FU_STATE_RSP. In any other case it will just answer by sending FU_STATE_RSP.

The content of FU_INIT_REQ is shown below.

Table 10-12 – Fields of FU_INIT_REQ

Field	Length	Description
Type	4 bits	0 = FU_INIT_REQ.
Version	2 bits	0 for this version of the protocol
PageSize	2 bits	0 for a PageSize=32 1 for a PageSize=64 2 for a PageSize=128 3 for a PageSize=192
ImageSize	32 bits	Size of the firmware upgrade image in bytes
CRC	32 bits	CRC of the firmware upgrade Image The input polynomial $M(x)$ is formed as a polynomial whose coefficients are bits of the data being checked (the first bit to check is the highest order coefficient and the last bit to check is the coefficient of order zero). The generator polynomial for the CRC is $G(x)=x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^8+x^7+x^5+x^4+x^2+x+1$. The remainder $R(x)$ is calculated as the remainder from the division of $M(x) \cdot x^{32}$ by $G(x)$. The coefficients of the remainder will then be the resulting CRC.

10.3.5.3.2 FU_EXEC_REQ

This packet is used by the base node to command a service node in "Complete" state to restart using the new firmware, once the complete image has been received by the service node. FU_EXEC_REQ specifies when the service node has to restart and how long the "safety" period shall be, as explained in clause 10.3.5.2.5. Additionally, FU_EXEC_REQ can be used in "Countdown" state to reset the restart and the safety timers.

Depending on the value of *RestartTimer*, a service node in "Complete" state may change either to "Countdown" or to "Upgrade" state. In any case, the service node answers with FU_STATE_RSP.

In "Countdown" state, the base node can reset *RestartTimer* and *SafetyTimer* with an FU_EXEC_REQ message (both timers must be specified in the message because both will be overwritten).

The content of this packet is described below.

Table 10-13 – Fields of FU_EXEC_REQ

Field	Length	Description
Type	4 bits	1 = FU_EXEC_REQ
Version	2 bits	0 for this version of the protocol
Reserved	2 bits	Reserved by ITU-T and set to 0.
RestartTimer	16 bits	0..65536 seconds; time before restarting with new FW
SafetyTimer	16 bits	0..65536 seconds; time to test the new FW. It starts when the "Upgrade" state is entered.

10.3.5.3.3 FU_CONFIRM_REQ

This packet is sent by the base node to a service node in "Upgrade" state to confirm the current FW. If the service node receives this message, it discards the old FW version and switches to "Idle" state. The service node answers with FU_STATE_RSP when receiving this message.

In any other state, the service node answers with FU_STATE_RSP without performing any additional actions.

This packet contains the fields described below.

Table 10-14 – Fields of FU_CONFIRM_REQ

Field	Length	Description
Type	4 bits	2 = FU_CONFIRM_REQ
Version	2 bits	0 for this version of the protocol
Reserved	2 bits	Reserved by ITU-T and set to 0.

10.3.5.3.4 FU_STATE_REQ

This packet is sent by the base node in order to get the firmware upgrade state of a service node. The service node will answer with FU_STATE_RSP.

This packet contains the fields described below.

Table 10-15 – Fields of FU_STATE_REQ

Field	Length	Description
Type	4 bits	3 = FU_STATE_REQ
Version	2 bits	0 for this version of the protocol
<i>Reserved</i>	2 bits	Reserved by ITU-T and set to 0.

10.3.5.3.5 FU_KILL_REQ

The base node sends this message to terminate the firmware upgrade process. A service node receiving this message will automatically switch to "Idle" state and optionally delete the downloaded data. The service node replies sending FU_STATE_RSP.

The content of this packet is described below.

Table 10-16 – Fields of FU_KILL_REQ

Field	Length	Description
Type	4 bits	4 = FU_KILL_REQ
Version	2 bits	0 for this version of the protocol
<i>Reserved</i>	2 bits	Reserved by ITU-T and set to 0.

10.3.5.3.6 FU_STATE_RSP

This packet is sent by the service node as an answer to FU_STATE_REQ, FU_KILL_REQ, FU_EXEC_REQ, FU_CONFIRM_REQ or FU_INIT_REQ messages received through the unicast connection. It is used to notify the firmware upgrade state in a service node.

Additionally, FU_STATE_RSP is used as a default response to all events that happen in states where they are not foreseen (e.g., FU_EXEC_REQ in "Receiving" state, FU_INIT_REQ in "Upgrade"...).

This packet contains the fields described below.

Table 10-17 – Fields of FU_STATE_RSP

Field	Length	Description
Type	4 bits	5 = FU_STATE_RSP
Version	2 bits	0 for this version of the protocol
Reserved	2 bits	Reserved by ITU-T and set to 0.
State	4 bits	0 for idle 1 for receiving 2 for complete 3 for countdown 4 for upgrade 5 to 15 reserved by ITU-T.
Reserved	4 bits	Reserved by ITU-T and set to 0.
CRC	32 bits	CRC as the one received in the CRC field of FU_INIT_REQ.
Received	32 bits	Number of received pages (this field should only be present if the state is "Receiving").

10.3.5.3.7 FU_DATA

This packet is sent by the base node to transfer a page of the firmware image to a service node. No answer is expected by the base node.

This packet contains the fields described below.

Table 10-18 – Fields of FU_DATA

Field	Length	Description
Type	4 bits	6 = FU_DATA
Version	2 bits	0 for this version of the protocol
<i>Reserved</i>	2 bits	Reserved by ITU-T and set to 0.
PageIndex	32 bits	Index of the page being transmitted
<i>Reserved</i>	8 bits	Padding byte for 16-bit devices. Set to 0 by default.
Data	<i>Variable</i>	Data of the page The length of this data is PageSize (32, 64, 128 or 192) bytes for every page, except the last one that will have the remaining bytes of the image.

10.3.5.3.8 FU_MISS_REQ

This packet is sent by the base node to a service node to request information about the pages that are still to be received.

If the service node is in "Receiving" state it will answer with an FU_MISS_BITMAP or FU_MISS_LIST message. If the service node is in any other state it will answer with an FU_STATE_RSP.

This packet contains the fields described below.

Table 10-19 – Fields of FU_MISS_REQ

Field	Length	Description
Type	4 bits	7 = FU_MISS_REQ
Version	2 bits	0 for this version of the protocol
<i>Reserved</i>	2 bits	Reserved by ITU-T and set to 0.
PageIndex	32 bits	Starting point to gather information about missing pages

10.3.5.3.9 FU_MISS_BITMAP

This packet is sent by the service node as an answer to an FU_MISS_REQ. It carries the information about the pages that are still to be received.

This packet will contain the fields described below.

Table 10-20 – Fields of FU_MISS_BITMAP

Field	Length	Description
Type	4 bits	8 = FU_MISS_BITMAP
Version	2 bits	0 for this version of the protocol
<i>Reserved</i>	2 bits	Reserved by ITU-T and set to 0.
PageIndex	32 bits	Page index of the page represented by the first bit of the bitmap. It should be the same as the <i>PageIndex</i> field in FU_MISS_REQ messages, or a posterior one. If it is posterior, it means that the pages in between are already received. In this case, if all pages after the <i>PageIndex</i> specified in FU_MISS_REQ have been received, the service node shall start looking from the beginning (<i>PageIndex</i> = 0).
Bitmap	<i>Variable</i>	This bitmap contains the information about the status of each page. The first bit (most significant bit of the first byte) represents the status of the page specified by <i>PageIndex</i> . The next bit represents the status of the <i>PageIndex</i> +1 and so on. A '1' represents that a page is missing, a '0' represents that the page is already received. After the bit that represents the last page in the image, it is allowed to overflow including bits that represent the missing status of the page with index zero. The maximum length of this field is <i>PageSize</i> bytes.

It is up to the service node to decide to send this type of packet or an FU_MISS_LIST message. It is usually more efficient to transmit this kind of packet when the number of missing packets is not very low. But it is up to the implementation to transmit one type of packet or the other. The base node should understand both.

10.3.5.3.10 FU_MISS_LIST

This packet is sent by the service node as an answer to an FU_MISS_REQ. It carries the information about the pages that are still to be received.

This packet will contain the fields described below.

Table 10-21 – Fields of FU_MISS_LIST

Field	Length	Description
Type	4 bits	9 = FU_MISS_LIST
Version	2 bits	0 for this version of the protocol
<i>Reserved</i>	2 bits	Reserved by ITU-T and set to 0.
PageIndexList	Variable	List of pages that are still to be received. Each page is represented by its <i>PageIndex</i> , coded as a 32 bit integer. These pages should be sorted in ascending order (low to high), being possible to overflow to the <i>PageIndex</i> equal to zero to continue from the beginning. The first page index should be the same as the <i>PageIndex</i> field in FU_MISS_REQ, or a posterior one. If it is posterior, it means that the pages in between are already received (by posterior it is allowed to overflow to the page index zero, to continue from the beginning). The maximum length of this field is <i>PageSize</i> bytes.

It is up to the service node to decide to transmit this packet type or an FU_MISS_BITMAP message. It is usually more efficient to transmit this kind of packet when the missing packets are very sparse, but it is implementation-dependent to transmit one type of packet or the other. The base node should understand both.

10.3.5.3.11 FU_INFO_REQ

This packet is sent by a base node to request information from a service node, such as manufacturer, device model, firmware version and other parameters specified by the manufacturer. The service node will answer with one or more FU_INFO_RSP packets.

This packet contains the fields described below.

Table 10-22 – Fields of FU_INFO_REQ

Field	Length	Description
Type	4 bits	10 = FU_INFO_REQ
Version	2 bits	0 for this version of the protocol
<i>Reserved</i>	2 bits	Reserved by ITU-T and set to 0.
InfoIdList	Variable	List of identifiers with the information to retrieve Each identifier is 1 byte long. The maximum length of this field is 32 bytes.

The following identifiers are defined:

Table 10-23 – InfoId possible values

InfoId	Name	Description
0	Manufacturer	Universal identifier of the manufacturer
1	Model	Model of the product working as service node
2	Firmware	Current firmware version being executed
128-255	<i>Manufacturer specific</i>	Range of values that are manufacturer specific

10.3.5.3.12 FU_INFO_RSP

This packet is sent by a service node as a response to an FU_INFO_REQ message from the base node. A service node may have to send more than one FU_INFO_RSP when replying to an information request by the base node.

This packet contains the fields described below.

Table 10-24 – Fields of FU_INFO_RSP

Field	Length	Description
Type	4 bits	11 = FU_INFO_RSP
Version	2 bits	0 for this version of the protocol
<i>Reserved</i>	2 bits	Reserved by ITU-T and set to 0.
InfoData	0-192 bytes	Data with the information requested by the base node. It may contain several entries (one for each requested identifier), each entry has a maximum size of 32 bytes. The maximum size of this field is 192 bytes (6 entries).

The InfoData field can contain several entries, the format of each entry is specified below.

Table 10-25 – Fields of each entry of InfoData in FU_INFO_RSP

Field	Length	Description
InfoId	8 bits	Identifier of the information as specified in 0
<i>Reserved</i>	3 bits	Reserved by ITU-T and set to 0.
Length	5 bits	Length of the data field (If length is 0 it means that the specified InfoId is not supported by the specified device).
Data	0-30 bytes	Data with the information provided by the service node. Its content may depend on the meaning of the InfoId field. No value may be longer than 30 bytes.

10.3.5.3.13 FU_CRC_REQ

FU_CRC_REQ is sent by the base node to command a service node to perform a CRC on the complete firmware image. The CRC on the complete FW image is mandatory. The CRC specified in FU_CRC_REQ.CRC is the same as in FU_INIT_REQ.

The service node replies with FU_CRC_RSP if it is in "Receiving" state, in any other case it replies with FU_STATE_RSP. The base node shall not send an FU_CRC_REQ if the image download is not complete (that is, the bitmap is not complete). Should the base node have an abnormal behaviour and send FU_CRC_REQ before completing the FW download, the service node would reply with FU_STATE_RSP.

Please note that in "Idle" state, the CRC field from FU_STATE_RSP will be a dummy (because no FU_INIT_REQ has been received yet). The base node will ignore this field if the service node is in "Idle" state.

This packet contains the fields described below.

Table 10-26 – Fields of FU_CRC_REQ

Field	Length	Description
Type	4 bits	12 = FU_CRC_REQ
Version	2 bits	0 for this version of the protocol
<i>Reserved</i>	2 bits	Reserved by ITU-T and set to 0.
SectionSize	32 bits	Size of the firmware upgrade image in bytes
CRC	32 bits	CRC of the firmware upgrade image The input polynomial $M(x)$ is formed as a polynomial whose coefficients are bits of the data being checked (the first bit to check is the highest order coefficient and the last bit to check is the coefficient of order zero). The generator polynomial for the CRC is $G(x)=x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^8+x^7+x^5+x^4+x^2+x+1$. The remainder $R(x)$ is calculated as the remainder from the division of $M(x) \cdot x^{32}$ by $G(x)$. The coefficients of the remainder will then be the resulting CRC.

10.3.5.3.14 FU_CRC_RSP

This packet is sent by the service node as a response to an FU_CRC_REQ message sent by the base node.

This packet contains the fields described below.

Table 10-27 – Fields of FU_CRC_RSP

Field	Length	Description
Type	4 bits	13 = FU_CRC_RSP
Version	2 bits	0 for this version of the protocol
CRC_Result	1 bit	Result of the CRC: "0" check failed "1" check OK.
<i>Reserved</i>	1 bit	Reserved by ITU-T and set to 0.

10.3.6 Examples

The figures below are examples of the traffic generated between the base node and the service node during the firmware upgrade process.

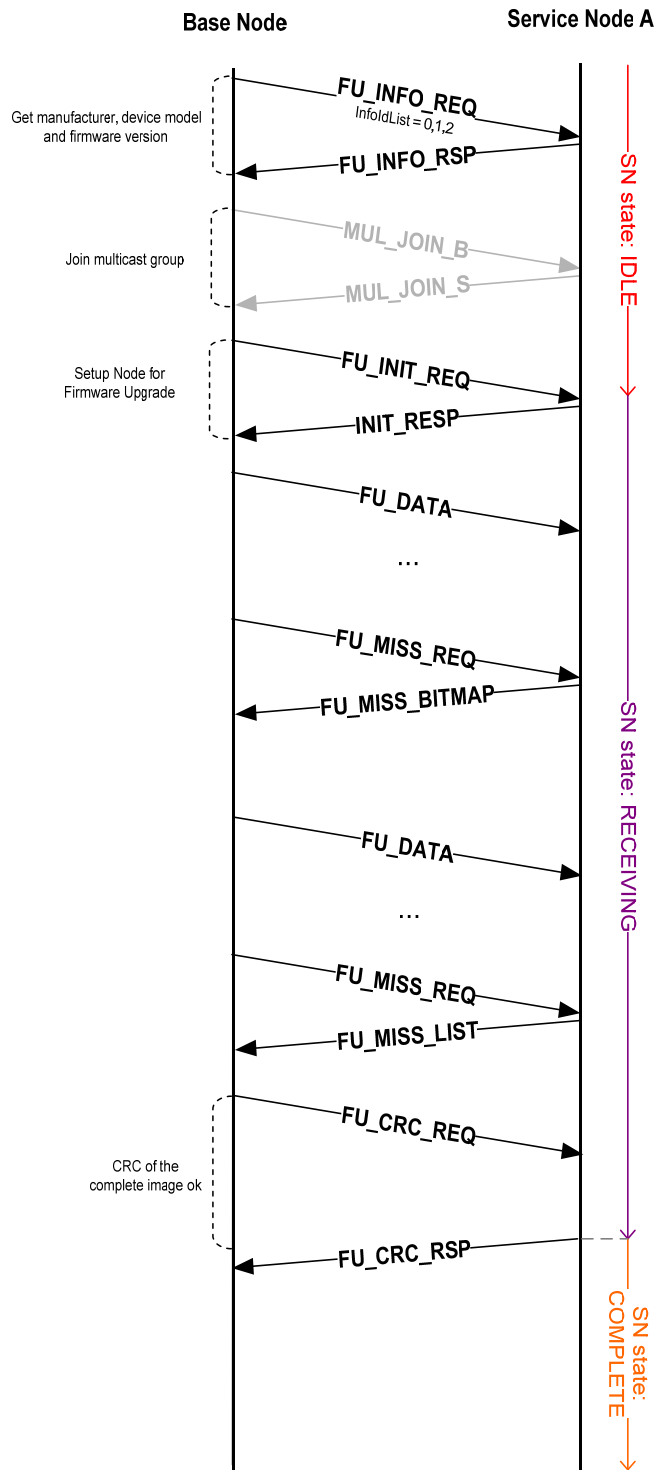


Figure 10-4 – Init service node and complete FW image

Figure 10-4 shows the initialization of the process, the FW download and the integrity check of the image. In the example above, the downloaded FW image is supposed to be complete before sending the last `FU_MISS_REQ`. The base node sends it to verify its bitmap. In this example, `FU_MISS_LIST` has an empty `PageIndexList` field, which means that the FW image is complete.

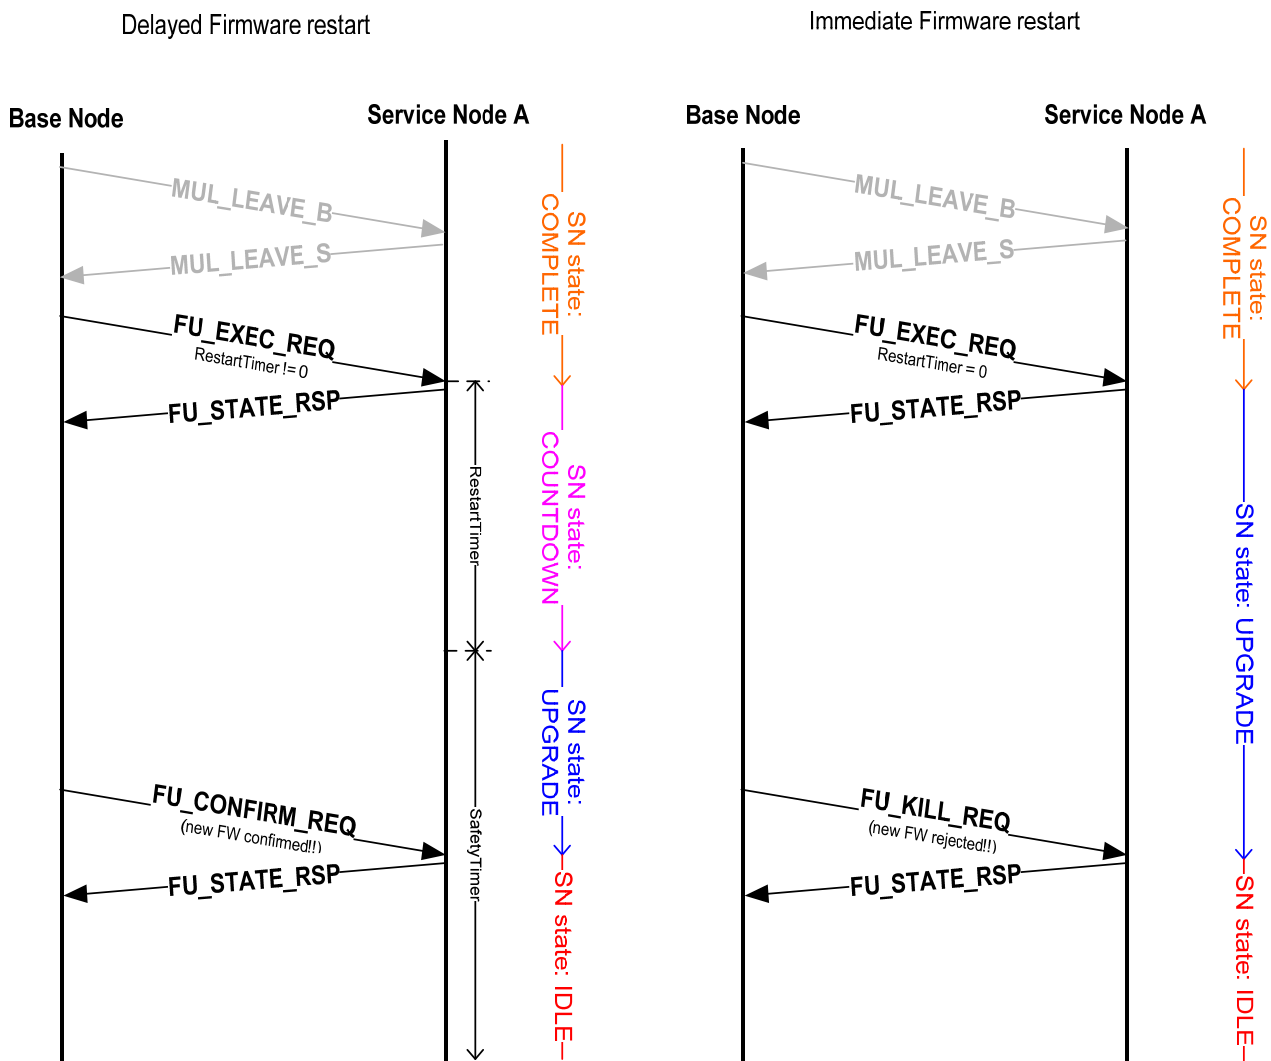


Figure 10-5 – Execute upgrade and confirm/reject new FW version

Figure 10-5 shows how to proceed after completing the FW download. The base node commands the service node to reboot either immediately ("Immediate Firmware Start", *RestartTimer* = 0) or after a defined period of time ("Delayed Firmware start", *RestartTimer* != 0). After reboot, the base node can either confirm the recently downloaded message sending an *FU_CONFIRM_REQ* or reject it (sending an *FU_KILL_REQ* or letting the safety period expire by doing nothing).

10.4 Management interface description

10.4.1 General

Management functions defined in earlier clauses shall be available over an abstract management interface specified in this clause. The management interface can be accessed over diverse media. Each physical media shall specify its own management plane communication profile over which management information is exchanged. It is mandatory for implementations to support the ITU-T G.9904 management plane communication profile. All other "management plane communication profiles" are optional and may be mandated by certain "application profiles" to use in specific cases.

This Recommendation describes two communication profiles, one of which is over this specification's NULL SSCS and the other over a serial link.

With these two communication profiles, it shall be possible to address the following use-cases:

1. Remote access of management interface over NULL SSCS. This shall enable base node's use as a supervisory gateway for all devices in a subnetwork.
2. Local access of management interface (over peripherals like RS232, USBSerial etc.) in a service node. Local access shall fulfil cases where a co-processor exists for supervisory control of the processor or when manual access is required over the local physical interface for maintenance.

Management data comprises of a 2 byte header followed by payload information corresponding to the type of information carried in the message. The header comprises of a 10 bit length field and a 6 bit message_id field.




Figure 10-6 – Management data frame format

Table 10-28 – Management data frame fields

Name	Length	Description
MGMT.LEN	10 bits	Length of payload data following the 2 byte header. LEN=0 implies there is no payload data following this header and the TYPE field contains all required information to perform appropriate action. NOTE – The length field may be redundant in some communication profiles (e.g., when transmitted over ITU-T G.9904), but it is required in others. Therefore for the sake of uniformity, it is always included in management data.
MGMT.TYPE	6 bits	Type of management information carried in corresponding data. Some message_ids have standard semantics which should be respected by all ITU-T G.9904 compliant devices while others are reserved for local use by vendors. 0x00 – Get PIB attribute query 0x01 – Get PIB attribute response 0x02 – Set PIB attribute command 0x03 – Reset all PIB statistics attributes 0x04 – Reboot destination device 0x05 – Firmware upgrade protocol message 0x06 – 0x0F: Reserved by ITU-T. Vendors should not use these values for local purposes. 0x10 – 0x3F: Reserved for vendor-specific use

10.4.2 Payload format of management information

10.4.2.1 Get PIB attribute query

This query is issued by a remote management entity that is interested in knowing the values of PIB attributes maintained on a compliant device with this Recommendation.

The payload may comprise of a query on either a single PIB attribute or multiple attributes. For reasons of efficiency queries on multiple PIB attributes may be aggregated in one single command.

Given that the length of a PIB attribute identifier is constant, the number of attributes requested in a single command is derived from the overall MGMT.LEN field in the header.

The format of payload information is shown in the following figure.




Figure 10-7 – Get PIB attribute query load

Fields of a GET request are summarized in the table below:

Table 10-29 – GET PIB attribute request fields

Name	Length	Description
PIB attribute ID	2 bytes	16 bit PIB attribute identifier
Index	1 byte	Index of entry to be returned for corresponding PIB attributeIDs. This field is only of relevance while returning PIB list attributes. Index = 0; if PIB attribute is not a list Index = 1 to 255; return list record at given index

10.4.2.2 Get PIB attribute response

This data is sent out from a compliant device of this Recommendation in response to a query of one or more PIB attributes. If a certain queried PIB attribute is not maintained on the device, it shall still respond to the query with a value field containing all-ones in the response.

The format of the payload is shown in the following figure.




Figure 10-8 – Get PIB attribute response.Payload

Fields of a GET request are summarized in table below:

Table 10-30 – GET PIB attribute response fields

Name	Length	Description
PIB attribute ID	2 bytes	16 bit PIB attribute identifier
Index	1 byte	Index of entry returned for corresponding PIB attributeIDs. This field is only of relevance while returning PIB list attributes. Index = 0; if PIB attribute is not a list Index = 1 to 255; returned list record is at given index
PIB attribute value	'a' bytes	Values of requested PIB attributes. In the case of a list attribute, the value shall comprise of an entire record corresponding to a given index of PIB attributes.

Table 10-30 – GET PIB attribute response fields

Name	Length	Description
Next	1 byte	Index of next entry returned for corresponding PIB AttributeIDs. This field is only of relevance while returning PIB list attributes. Next = 0; if PIB attribute is not a list or if no records follow the one being returned for a list PIB attribute, i.e., given record is last entry in list. Next = 1 to 255; index of next record in list maintained for given PIB attribute.

Response to a PIB attribute query can span across several MAC GPDU's. This shall always be the case when an aggregated (comprising of several PIB attributes) PIB query's response is longer than the maximum segment size is allowed to be carried over the NULL SCSS.

10.4.2.3 Set PIB attribute

This management data shall be used to set specific PIB attributes. Such management payload comprises of a 2 byte PIB attribute identifier, followed by the relevant length of PIB attribute information corresponding to that identifier. For reasons of efficiency, it shall be possible to aggregate a SET command on several PIB attributes in one GPDU. The format of such an aggregated payload is shown in Figure 10-9 below:




Figure 10-9 – Set PIB attribute query payload

For cases where the corresponding PIB attribute is only a trigger (all ACTION PIB attributes), there shall be no associated value and the request data format shall be as shown below:




Figure 10-10 – Set PIB attribute response payload

It is assumed that the management entity sending out this information has already determined if the corresponding attributes are supported at the target device. This may be achieved by a previous query and its response.

10.4.2.4 Reset statistics

This command has optional payload. In case there is no associated payload, the receiving device shall reset all of its PIB statistical attributes.

For cases when a remote management entity only intends to perform the reset of selective PIB statistical attributes, the payload shall contain a list of attributes that need to be reset. The format shall be the same as shown in clause 10.4.2.1.

Since there is no confirmation message going back from the device complying with this Recommendation, the management entity needs to send a follow-up PIB attribute query, in case it wants to confirm the successful completion of appropriate action.

10.4.2.5 Reboot device

There is no corresponding payload associated with this command. The command is complete in itself. The receiving compliant device with this Recommendation shall reboot itself upon receipt of this message.

It is mandatory for all implementations compliant with this Recommendation to support this command and its corresponding action.

10.4.2.6 Firmware upgrade

The payload in this case shall comprise of firmware upgrade commands and responses described in clause 10.3 of the Recommendation.

10.4.3 NULL SCS communication profile

This communication profile enables the exchange of management information described in previous clauses over the NULL SCS.

The management entities at both transmitting and receiving ends are applications making use of the NULL SCS enumerated in clause 9.3 of this Recommendation. Data is therefore exchanged as MAC generic PDUs.

10.4.4 Serial communication profile

10.4.4.1 Physical layer

The PHY layer may be any serial link (e.g., RS232, USB Serial). The serial link is required to work 8N1 configuration at one of the following data rates:

9600 bps, 19200 bps, 38400 bps, 57600 bps

10.4.4.2 Data encapsulation for management messages

In order to ensure robustness, the stream of data is encapsulated in HDLC-type frames which include a 2 byte header and a 4 byte CRC. All data is encapsulated between a starting flag-byte 0x7E and an ending flag-byte 0x7E as shown in Figure 10-11 below:




Figure 10-11 – Data encapsulation for management messages

If any of the intermediate data characters has the value 0x7E, it is preceded by an escape byte 0x7D, followed by a byte derived from XORing the original character with byte 0x20. The same is done if there is a 0x7D within the character stream. An example of such a case is shown here:

```
Msg to Tx:      0x01 0x02 0x7E  0x03 0x04 0x7D  0x05 0x06
Actual Tx sequence: 0x01 0x02 0x7D 0x5E 0x03 0x04 0x7D 0x5D 0x05 0x06
                Escape      Escape
                sequence     sequence
```

The 32 bit CRC at the end of the frame covers both '*Header*' and '*Payload*' fields. The CRC is calculated over the original data to be transmitted, i.e., before byte stuffing of the escape sequences described above is performed. The CRC calculation is as follows:

The input polynomial $M(x)$ is formed as a polynomial whose coefficients are bits of the data being checked (the first bit to check is the highest order coefficient and the last bit to check is the coefficient of order zero). The generator polynomial for the CRC is $G(x)=x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^8+x^7+x^5+x^4+x^2+x+1$. The remainder $R(x)$ is calculated as the remainder from the division of $M(x) \cdot x^{32}$ by $G(x)$. The coefficients of the remainder will then be the resulting CRC.

10.5 List of mandatory PIB attributes

10.5.1 General

PIB attributes listed in this clause shall be supported by all implementations. PIB attributes that are not listed in this clause are optional and vendors may implement them at their choice. In addition to the PIB attributes, the management command to reboot a certain device (as specified in clause 10.4.2.5) shall also be universally supported.

10.5.2 Mandatory PIB attributes common to all device types

10.5.2.1 PHY PIB attribute

(See Table 10-1.)

Table 10-31 – PHY PIB common mandatory attributes

Attribute name	Id
phyStatsRxTotalCount	0x00A4
phyStatsBlkAvgEvm	0x00A5
phyEmaSmoothing	0x00A8

10.5.2.2 MAC PIB attributes

(See Tables 10-3, 10-5 and 10-7.)

Table 10-32 – MAC PIB common mandatory attributes

Attribute name	Id
macEMASmoothing	0x0019
macCapabilities	0x002C
macListPhyComm	0x0057

10.5.2.3 Application PIB attributes

(See Table 10-9.)

Table 10-33 – Applications PIB common mandatory attributes

Attribute name	Id
appFwVersion	0x0075
appVendorId	0x0076
appProductId	0x0077

10.5.3 Mandatory base node attributes

10.5.3.1 MAC PIB attributes

(See Tables 10-3 and 10-7.)

Table 10-34 – MAC PIB base node mandatory attributes

Attribute name	Id
macBeaconsPerFrame	0x0013
macListRegDevices	0x0050
macListActiveConn	0x0051

10.5.4 Mandatory service node attributes

10.5.4.1 MAC PIB attributes

(See Tables 10-5, 10-7 and 10-8.)

Table 10-35 – MAC PIB service node mandatory attributes

Attribute name	Id
macLNID	0x0020
macLSID	0x0021
macSID	0x0022
macSNA	0x0023
macState	0x0024
macSCPLength	0x0025
macNodeHierarchyLevel	0x0026
macBeaconSlotCount	0x0027
macBeaconRxSlot	0x0028
macBeaconTxSlot	0x0029
macBeaconRxFrequency	0x002A
macBeaconTxFrequency	0x002B
macListSwitchTable	0x0053
macListAvailableSwitches	0x0056
MACActionTxData	0x0060
MACActionConnClose	0x0061
MACActionRegReject	0x0062
MACActionProReject	0x0063
MACActionUnregister	0x0064

10.5.4.2 Application PIB attributes

(See Table 10-10.)

Table 10-36 – APP PIB service node mandatory attributes

Attribute name	Id
appFwdlRunning	0x0070
appFwdlRxPktCount	0x0071

Annex A

EVM and SNR calculation

(This annex forms an integral part of this Recommendation.)

This annex describes calculations of the EVM by a reference receiver, assuming accurate synchronization and FFT window placement.

Let

$\{r_k^i; k = 1, 2, \dots, 97\}$ denotes that the FFT output for symbol i and k are the frequency tones.

$\Delta b_k \in \{0, 1, \dots, M-1\}$ represents the decision on the received information symbol coded in the phase increment.

$M = 2, 4$ or 8 in the case of DBPSK, DQPSK or D8PSK, respectively.

The EVM definition is then given by;

$$EVM = \frac{\sum_{i=1}^L \sum_{k=2}^{97} \left[\text{abs} \left(r_k^i - r_{k-1}^i e^{-\left(\frac{j2\pi}{M}\right) \Delta b_{k-1}} \right) \right]^2}{\sum_{i=1}^L \sum_{k=2}^{97} \left[\text{abs}(r_k^i) \right]^2}$$

In the above, $\text{abs}(\cdot)$ refers to the magnitude of a complex number. L is the number of OFDM symbols in the payload of the most recently received PPDU, over which the EVM is calculated.

The SNR is then defined as the reciprocal of the EVM above.

Annex B

MAC layer constants

(This annex forms an integral part of this Recommendation.)

This clause defines all the MAC layer constants.

Table B.1 – Table of MAC constants

Constant	Value	Description
MACBeaconLength	4 symbols	Length of beacon in symbols
MACMinSCPLength	64 symbols	Minimum length of SCP
MACPriorityLevels	4	Number of levels of priority supported by the system
MACCFPMaxAlloc	32	Maximum contention-free periods that may be allocated within a frame
MACFrameLength	276 symbols	Length of a frame in number of symbols
MACRandSeqChgTime	32767 seconds (approx 9 hours)	Maximum duration of time after which the base node should circulate a new random sequence to the subnetwork for encryption functions.
MACMaxPRNIgnore	3	Maximum number of promotion-needed messages a terminal can ignore.
$N_{\text{miss-beacon}}$	5	Number of times a service node does not receive an expected beacon before considering its switch node as unavailable.
ARQMaxTxCount	5	Maximum transmission count before declaring a permanent failure.
ARQCongClrTime	2 sec	When the receiver has indicated congestion, this time must be waited before retransmitting the data.
ARQMaxCongInd	7	After ARQMaxCongInd consecutive transmissions which failed due to congestion, the connection should be declared permanently dead.
ARQMaxAckHoldTime	7 sec	Time the receiver may delay sending an ACK in order to allow consolidated ACKs or piggyback the ACK with a data packet.

Annex C

Convergence layer constants

(This annex forms an integral part of this Recommendation.)

The following TYPE values are defined for use by the convergence layers from clause 9.

Table C.1 – TYPE value assignments

TYPE symbolic name	Value
TYPE_CL_IPv4_AR	1
TYPE_CL_IPv4_UNICAST	2
TYPE_CL_432	3
TYPE_CL_MGMT	4
TYPE_CL_IPv6_AR	5
TYPE_CL_IPv6_UNICAST	6

The following LCID values apply for broadcast connections defined by the convergence layers from clause 9.

Table C.2 – LCID value assignments

LCID Symbolic Name	Value	MAC scope
LCI_CL_IPv4_BROADCAST	1	Broadcast.
LCI_CL_432_BROADCAST	2	Broadcast.

The following result values are defined for convergence layer primitives.

Table C.3 – Result values for convergence layer primitives

Result	Description
Success = 0	The SSCS service was successfully performed.
Reject = 1	The SSCS service failed because it was rejected by the base node.
Timeout = 2	A timed out occurs during the SSCS service processing.
Not Registered = 6	The service node is not currently registered to a subnetwork.

Appendix I

Examples of CRC

(This appendix does not form an integral part of this Recommendation.)

The table below gives the CRCs calculated for several specified strings.

Table I.1 – Examples of CRCs calculated for various ASCII strings

String	CRC-8
'T'	0xab
"THE"	0xa0
0x03, 0x73	0x61
0x01, 0x3f	0xa8
"123456789"	0xf4

Appendix II

Interleaving matrices

(This appendix does not form an integral part of this Recommendation.)

Header interleaving matrix:

12	11	10	9	8	7	6	5	4	3	2	1
24	23	22	21	20	19	18	17	16	15	14	13
36	35	34	33	32	31	30	29	28	27	26	25
48	47	46	45	44	43	42	41	40	39	38	37
60	59	58	57	56	55	54	53	52	51	50	49
72	71	70	69	68	67	66	65	64	63	62	61
84	83	82	81	80	79	78	77	76	75	74	73

DBPSK (FEC ON) interleaving matrix:

12	11	10	9	8	7	6	5	4	3	2	1
24	23	22	21	20	19	18	17	16	15	14	13
36	35	34	33	32	31	30	29	28	27	26	25
48	47	46	45	44	43	42	41	40	39	38	37
60	59	58	57	56	55	54	53	52	51	50	49
72	71	70	69	68	67	66	65	64	63	62	61
84	83	82	81	80	79	78	77	76	75	74	73
96	95	94	93	92	91	90	89	88	87	86	85

DQPSK (FEC ON) interleaving matrix:

12	11	10	9	8	7	6	5	4	3	2	1
24	23	22	21	20	19	18	17	16	15	14	13
36	35	34	33	32	31	30	29	28	27	26	25
48	47	46	45	44	43	42	41	40	39	38	37
60	59	58	57	56	55	54	53	52	51	50	49
72	71	70	69	68	67	66	65	64	63	62	61
84	83	82	81	80	79	78	77	76	75	74	73
96	95	94	93	92	91	90	89	88	87	86	85
108	107	106	105	104	103	102	101	100	99	98	97
120	119	118	117	116	115	114	113	112	111	110	109
132	131	130	129	128	127	126	125	124	123	122	121
144	143	142	141	140	139	138	137	136	135	134	133
156	155	154	153	152	151	150	149	148	147	146	145
168	167	166	165	164	163	162	161	160	159	158	157
180	179	178	177	176	175	174	173	172	171	170	169
192	191	190	189	188	187	186	185	184	183	182	181

D8PSK (FEC ON) interleaving matrix:

18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19
54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37
72	71	70	69	68	67	66	65	64	63	62	61	60	59	58	57	56	55
90	89	88	87	86	85	84	83	82	81	80	79	78	77	76	75	74	73
108	107	106	105	104	103	102	101	100	99	98	97	96	95	94	93	92	91
126	125	124	123	122	121	120	119	118	117	116	115	114	113	112	111	110	109
144	143	142	141	140	139	138	137	136	135	134	133	132	131	130	129	128	127
162	161	160	159	158	157	156	155	154	153	152	151	150	149	148	147	146	145
180	179	178	177	176	175	174	173	172	171	170	169	168	167	166	165	164	163
198	197	196	195	194	193	192	191	190	189	188	187	186	185	184	183	182	181
216	215	214	213	212	211	210	209	208	207	206	205	204	203	202	201	200	199
234	233	232	231	230	229	228	227	226	225	224	223	222	221	220	219	218	217
252	251	250	249	248	247	246	245	244	243	242	241	240	239	238	237	236	235
270	269	268	267	266	265	264	263	262	261	260	259	258	257	256	255	254	253
288	287	286	285	284	283	282	281	280	279	278	277	276	275	274	273	272	271

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Terminals and subjective and objective assessment methods
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects and next-generation networks
Series Z	Languages and general software aspects for telecommunication systems