

Specification for *PowerLine Intelligent Metering Evolution*



Prepared by the PRIME Alliance Technical Working Group

This document is an approved specification. As such, it is subject to change. Prior to the full or partial adoption of this document by any standards development organization, permission must be obtained from the PRIME Alliance.

Abstract:

This is a complete specification for a new OFDM-based power line communication system for the provision of all kinds of Smart Grid services over electricity distribution networks. Both PHY and MAC layers according to IEEE conventions, plus a Convergence layer, are described in the Specification.

Content Table

Content Table	2
List of Figures	9
List of Tables	16
1 Introduction	23
1.1 Scope	23
1.2 Overview.....	23
1.3 Normative references.....	23
1.4 Document conventions	26
1.5 Definitions	27
1.6 Abbreviations and Acronyms.....	28
2 General Description	34
2.1 Introduction.....	34
2.2 General description of the architecture	34
3 Physical layer	35
3.1 Introduction.....	35
3.2 Overview.....	35
3.2.1 General	35
3.2.2 Note about backwards compatibility with PRIME v1.3.6.....	37
3.3 PHY parameters.....	37
3.4 Preamble, header and payload structure.....	39
3.4.1 Preamble.....	39
3.4.2 Pilot structure	42
3.4.3 Header and Payload.....	45
3.5 Convolutional encoder	50
3.6 Scrambler.....	50

3.7	Repeater	51
3.8	Interleaver	52
3.9	Modulation	53
3.10	Electrical specification of the transmitter	55
3.10.1	General	55
3.10.2	Transmit PSD.....	55
3.10.3	Error Vector Magnitude (EVM).....	57
3.10.4	Conducted disturbance limits.....	57
3.11	PHY service specification	57
3.11.1	General	57
3.11.2	PHY Data plane primitives	58
3.11.3	PHY Control plane primitives.....	62
3.11.4	PHY Management primitives.....	69
4	MAC layer	75
4.1	Overview.....	75
4.2	Addressing	76
4.2.1	General	76
4.2.2	Example of address resolution	77
4.2.3	Broadcast and multicast addressing.....	79
4.3	MAC functional description.....	79
4.3.1	Service Node start-up	79
4.3.2	Starting and maintaining Subnetworks	80
4.3.3	Channel Access	81
4.3.4	Tracking switches and peers.....	87
4.3.5	Switching	90
4.3.6	Direct connections.....	92

4.3.7	Packet aggregation	97
4.3.8	Security	98
4.4	MAC PDU format	111
4.4.1	General	111
4.4.2	Generic MAC PDU	111
4.4.3	Promotion Needed PDU	143
4.4.4	Beacon PDU	145
4.5	MAC Service Access Point.....	148
4.5.1	General	148
4.5.2	Service Node and Base Node signalling primitives.....	150
4.5.3	Base Node signalling primitives.....	159
4.5.4	Service and Base Nodes data primitives.....	160
4.5.5	MAC Layer Management Entity SAPs.....	161
4.6	MAC procedures.....	171
4.6.1	Registration process	171
4.6.2	Unregistration process	173
4.6.3	Promotion process.....	174
4.6.4	Demotion process.....	178
4.6.5	Keep-Alive process.....	179
4.6.6	Connection establishment.....	184
4.6.7	Multicast group management	186
4.6.8	Robustness Management.....	193
4.6.9	Channel allocation	195
4.7	Automatic Repeat Request (ARQ)	197
4.7.1	General	197
4.7.2	Initial negotiation	197

4.7.3	ARQ mechanism	197
4.7.4	ARQ packets switching	201
4.8	Time Reference.....	201
4.9	Backward Compatibility with PRIME 1.3.6	202
4.9.1	Frame Structure and Channel Access	203
4.9.2	Switching	204
4.9.3	PDU Frame Formats.....	204
5	Convergence layer	206
5.1	Overview.....	206
5.2	Common Part Convergence Sublayer (CPCS)	206
5.2.1	General	206
5.2.2	Segmentation and Reassembly (SAR).....	206
5.3	NULL Service-Specific Convergence Sublayer (NULL SSCS)	209
5.3.1	Overview.....	209
5.3.2	Primitives	209
5.4	IPv4 Service-Specific Convergence Sublayer (IPv4 SSCS)	210
5.4.1	Overview.....	210
5.4.2	Address resolution.....	211
5.4.3	IPv4 packet transfer.....	213
5.4.4	Segmentation and reassembly	214
5.4.5	Header compression.....	214
5.4.6	Quality of Service mapping.....	215
5.4.7	Packet formats and connection data.....	215
5.4.8	Service Access Point	222
5.5	IEC 61334-4-32 Service-Specific Convergence Sublayer (IEC 61334-4-32 SSCS)	228
5.5.1	General	228

5.5.2	Overview	228
5.5.3	Address allocation and connection establishment	228
5.5.4	Connection establishment data format	230
5.5.5	Packet format	231
5.5.6	Service Access Point	231
5.6	IPv6 Service-Specific Convergence Sublayer (IPv6 SCS)	233
5.6.1	Overview	233
5.6.2	IPv6 Convergence layer	234
5.6.3	IPv6 Address Configuration	235
5.6.4	IPv6 Packet Transfer	237
5.6.5	Segmentation and reassembly	240
5.6.6	Compression	240
5.6.7	Quality of Service Mapping	241
5.6.8	Packet formats and connection data.....	242
5.6.9	Service access point.....	248
6	Management plane	254
6.1	Introduction.....	254
6.2	Node management.....	255
6.2.1	General	255
6.2.2	PHY PIB attributes.....	255
6.2.3	MAC PIB attributes	258
6.2.4	Application PIB attributes.....	281
6.3	Firmware upgrade	283
6.3.1	General	283
6.3.2	Requirements and features	283
6.3.3	General Description.....	283

6.3.4	Firmware upgrade PIB attributes	286
6.3.5	State machine	286
6.3.6	Examples	304
6.4	Management interface description	306
6.4.1	General	306
6.4.2	Payload format of management information	308
6.4.3	NULL SSCS communication profile	313
6.4.4	Serial communication profile	313
6.4.5	TCP communication profile	314
6.5	List of mandatory PIB attributes	314
6.5.1	General	314
6.5.2	Mandatory PIB attributes common to all device types	314
6.5.3	Mandatory Base Node attributes	316
6.5.4	Mandatory Service Node attributes	316
Annex A (informative)	Examples of CRC	319
Annex B (normative)	EVM calculation	320
Annex C (informative)	Interleaving matrixes ($N_{CH} = 1$)	321
Annex D (normative)	MAC layer constants	324
Annex E (normative)	Convergence layer constants	326
Annex F (normative)	Profiles	328
F.1	Smart Metering Profile	328
Annex G (informative)	List of frequencies used	330
Annex H (informative)	Informative	343
H.1	Data exchange between to IP communication peers	343
H.2	Joining a multicast group	345
Annex I (informative)	ARQ algorithm	347

Annex J (normative) PHY backwards compatibility mechanism with PRIME v1.3.6	348
Annex K (normative) MAC Backward Compatibility PDUs and Procedures	352
K.1 MAC PDU format	352
K.1.1 Generic MAC PDU	352
K.1.2 Compatibility Beacon PDU (CBCN)	365
K.2 MAC procedures.....	368
K.2.1 Registration process	368
K.2.2 Unregistering process.....	369
K.2.3 Promotion process.....	369
K.2.4 Demotion process.....	372
K.2.5 Keep-Alive process.....	373
K.2.6 Connection management	374
K.2.7 Multicast group management	374
K.2.8 Robustness Management.....	374
K.2.9 Channel allocation and deallocation	374
Annex L (Informative) Type A, Type B PHY frames and Robust modes.....	375
List of authors (by alphabetical order).....	376

List of Figures

Figure 1 - Reference model of protocol layers used in the PRIME specification	34
Figure 2 - Overview of PPDU processing	36
Figure 3 - PHY frame of Type A format.....	36
Figure 4 - PHY frame of Type B format.....	36
Figure 5 - Example of the preamble structure when three channels are used.....	41
Figure 6 - Preamble Type B structure.....	41
Figure 7 - Example of each of the preamble symbol structure when three channels are used	42
Figure 8 - Pilot and data subcarrier frequency allocation inside the header (eight active channels case).....	43
Figure 9 - LFSR for use in Pilot sequence generation	43
Figure 10 - PHY frame of Type A, pilot and data subcarrier allocation (eight active channels case).....	44
Figure 11 - PHY frame of Type B, pilot and data subcarrier allocation (eight active channels case)	45
Figure 12 - PRIME PPDU of Type A: header and payload (bits transmitted before encoding)	46
Figure 13 - PRIME PPDU of Type B: header and payload (bits transmitted before encoding).....	48
Figure 14 - Convolutional encoder	50
Figure 15 - LFSR for use in the scrambler block.....	51
Figure 16 - Example of repeater block using a shift value = 2.....	51
Figure 17 - DBPSK, DQPSK and D8PSK mapping	53
Figure 18 - Subcarrier Mapping	54
Figure 19 - Measurement set up (single-phase).....	55
Figure 20 - Measurement set up (three-phase)	55
Figure 21 - Artificial mains network	56
Figure 22 - EVM meter (block diagram)	57
Figure 23 - Overview of PHY primitives.....	58
Figure 24 - Overview of PHY Control Plane Primitives	62
Figure 25 - Service Node states	75

Figure 26 - Addressing Structure	77
Figure 27 - Example of address resolution: phase 1.....	78
Figure 28 - Example of address resolution: phase 2.....	78
Figure 29 - Example of address resolution: phase 3.....	78
Figure 30 - Example of address resolution: phase 4.....	79
Figure 31 - Structure of a MAC Frame.....	82
Figure 32 - Flow chart for CSMA-CA algorithm	86
Figure 33 -Switching tables examples	88
Figure 34 - Filling example for the switching table for Switch of Level 0.....	89
Figure 35 - Directed Connection to an unknown Service Node	93
Figure 36 - Example of direct connection: connection establishment to a known Service Node.....	95
Figure 37 - Release of a direct connection	96
Figure 38 - Nonce structure.....	101
Figure 39 - Counter handling in ALV and request/response messages.....	103
Figure 40 - Counter handling in direct connections and request messages	104
Figure 41 - REG Nonce structure	105
Figure 42 - Key derivation hierarchy	105
Figure 43 - Security profile 1 and 2 encryption algorithm (authentication only).....	109
Figure 44 - Security profile 1 and 2 encryption algorithm (authentication and encryption).....	110
Figure 45 - Generic MAC PDU format.....	111
Figure 46 - Generic MAC header	111
Figure 47 - Packet structure	112
Figure 48 - Packet Header	112
Figure 49 - PKT.CID structure.....	113
Figure 50 - Security subheader.....	115
Figure 51 - Two transactions without requiring retransmits	118

Figure 52 - Transaction with packet loss requiring retransmits	119
Figure 53 - Duplicate packet detection and elimination	119
Figure 54 - REG control packet structure	125
Figure 55 - CON control packet structure.....	127
Figure 56 - PRO_REQ_S control packet structure	130
Figure 57 - PRO control packet structure	130
Figure 58 - FRA control packet structure.....	134
Figure 59 - CFP control packet structure	134
Figure 60 - ALV_RSP_S / ALV_REQ_B Control packet structure	136
Figure 61 - ALV_ACK_B/ALV_ACK_S control packet structure	136
Figure 62 - ALV_RSP_ACK Control packet structure.....	136
Figure 63 - MUL control packet structure	139
Figure 64 - SEC control packet structure	142
Figure 65 - Promotion Need MAC PDU	143
Figure 66 - Beacon PDU structure	145
Figure 67 - Establishment of a Connection.....	149
Figure 68 - Failed establishment of a Connection	149
Figure 69 - Release of a Connection	149
Figure 70- Transfer of Data.....	149
Figure 71 - Registration process accepted	172
Figure 72 - Registration process rejected.....	172
Figure 73 - Unregistration process initiated by a Terminal Node	174
Figure 74 - Unregistration process initiated by the Base Node.....	174
Figure 75 - Promotion process initiated by a Service Node	175
Figure 76 - Promotion process rejected by the Base Node.....	175
Figure 77 - Promotion process initiated by the Base Node.....	175

Figure 78 - Promotion process rejected by a Service Node	176
Figure 79 - BCN modulation change request initiated by the Switch.	176
Figure 80 - BCN modulation change request initiated by the Switch and rejected by the Base Node.....	177
Figure 81 - BCN modulation change request initiated by the Base Node.....	177
Figure 82 - BCN modulation change request initiated by the Base Node and rejected by the Switch.....	177
Figure 83 - Demotion process initiated by a Service Node	179
Figure 84 - Demotion process initiated by the Base Node.....	179
Figure 85 - Successful ALV procedure	180
Figure 86 - Failed ALV procedure	181
Figure 87 - Failed ALV procedure (uplink)	181
Figure 88 - ALV.REP_U Example (ALV.MIN_LEVEL = 2)	182
Figure 89 - ALV.REP_D Example (ALV.MIN_LEVEL = 2).....	183
Figure 90 - Connection establishment initiated by a Service Node	184
Figure 91 - Connection establishment rejected by the Base Node.....	184
Figure 92 - Connection establishment initiated by the Base Node.....	185
Figure 93 - Connection establishment rejected by a Service Node	185
Figure 94 - Disconnection initiated by a Service Node.....	186
Figure 95 - Disconnection initiated by the Base Node	186
Figure 96 - Successful group join initiated by Base Node.....	187
Figure 97 - Failed group join initiated by Base Node	188
Figure 98 - Successful group join initiated by Service Node.....	189
Figure 99 - Failed group join initiated by Service Node.....	190
Figure 100 - Leave initiated by the Base Node.....	191
Figure 101 - Leave initiated by the Service Node	191
Figure 102 - Multicast Switching Tracking for Group Join.....	192
Figure 103 - Multicast Switching Tracking for Group Leave.....	193

Figure 104 - Successful allocation of CFP period	196
Figure 105 - Successful channel de-allocation sequence	196
Figure 106 - Deallocation of channel to one device results in the change of CFP allocated to another	197
Figure 107 - ARQ subheader only with the packet id.....	198
Figure 108 - ARQ subheader with ARQ.INFO	198
Figure 109 - ARQ.ACK byte fields.....	198
Figure 110 - ARQ.WIN byte fields	198
Figure 111 - ARQ.NACK byte fields	198
Figure 112 - Example of an ARQ subheader with all the fields present.....	199
Figure 113 - Stop and wait ARQ subheader with only packet ID	200
Figure 114 - Stop and wait ARQ subheader with an ACK	201
Figure 115 - Stop and wait ARQ subheader without data and with an ACK	201
Figure 116 - CBCN Frame format for backwards compatibility mode.....	203
Figure 117 - Robust beacon allocation in 1.4 BC	203
Figure 118 - Structure of the Convergence layer	206
Figure 119 - Segmentation and Reassembly Headers	207
Figure 120 - IPv4 SSCS connection example.....	211
Figure 121 - IPv6 SSCS connection example.....	234
Figure 122 - Management plane. Introduction.	254
Figure 123 - Restarting de nodes and running the new firmware	285
Figure 124 - Signed firmware diagram	285
Figure 125 - Firmware Upgrade mechanism, state diagram	290
Figure 126 - Init Service Node and complete FW image	305
Figure 127 - Execute upgrade and confirm/reject new FW version.....	306
Figure 128 - Management data frame	307
Figure 129 - Get PIB Attribute query. Payload	308

Figure 130 - Get PIB Attribute response. Payload	309
Figure 131 - Set PIB attribute. Aggregated payload	310
Figure 132 - Set PIB Attribute. ACTION PIB attributes	310
Figure 133 - Enhanced PIB query format.....	311
Figure 134 - Iterator short format	311
Figure 135 - Iterator long format.....	311
Figure 136 - Enhanced PIB response format	312
Figure 137 - Data encapsulations for management messages.....	313
Figure 138 - Backwards Compatible PHY frame	348
Figure 139 - BC frame predefined content.....	349
Figure 140 - PHY BC frame detected by v1.3.6 devices.....	350
Figure 141 - BC PHY frame detected by v1.4 devices.....	350
Figure 142 - v1.3.6 frame received by a v1.4 node	351
Figure 143 - BC PHY frame in noisy environment.....	351
Figure 144 - Compatibility Packet Header	352
Figure 145 - CREG control packet structure	354
Figure 146 - CPRO_REQ_S control packet structure	359
Figure 147 - CPRO control packet structure	359
Figure 148 - CBSI control packet structure.....	362
Figure 149 - CFRA control packet structure.....	363
Figure 150 - CALV Control packet structure	364
Figure 151 - Beacon PDU structure	365
Figure 152 - Registration process accepted	368
Figure 153 - Registration process rejected.....	369
Figure 154 - Promotion process initiated by a Service Node	370
Figure 155 - Promotion process rejected by the Base Node.....	371

Figure 156 - Promotion process initiated by the Base Node.....	371
Figure 157 - Promotion process rejected by a Service Node	371
Figure 158 - Double switching BSI message exchange	372
Figure 159 - Demotion process initiated by a Service Node	373
Figure 160 - Demotion process initiated by the Base Node.....	373

List of Tables

Table 1 - Frequency and Timing Parameters of the PRIME PHY	37
Table 2 - PHY Payload Parameters	38
Table 3 - PHY Header Parameters	39
Table 4 - Roll-off region length for all N_{CH} values	40
Table 5 - Roll-off region length for all N_{CH} values	42
Table 6 - Length in bits of MPDU1 and PAD_H fields in the PHY frame header of Type A for all possible values of N_{CH}	47
Table 7 - Length in bits of MPDU1 and PAD_H fields in the PHY frame header of Type B for all possible values of N_{CH}	49
Table 8 - Shift values for the Robust modes.....	52
Table 9 - Fields associated with PHY Control Plane Primitives.....	62
Table 10 - Values of the status parameter in PLME_GET.confirm primitive	74
Table 11 - Broadcast and multicast address.....	79
Table 12 - Direct connection example: Node B's Direct switching table	94
Table 13 - Values of <i>MACUpdateKeysTime</i> for different number of channels	100
Table 14 -Nonce SID/LNID Derivation	101
Table 15 - Encryption/Authentication by PDU Types.....	106
Table 16 - Generic MAC header fields	111
Table 17 - Packet header fields	113
Table 18 - Security subheader fields	115
Table 19 - MAC control packet types	116
Table 20 - REG control packet fields.....	120
Table 21 - REG control packet types.....	126
Table 22 - CON control packet fields	127
Table 23 - CON control packet types.....	129

Table 24 - PRO control packet fields.....	131
Table 25 - PRO control packet types	133
Table 26 - FRA control packet fields.....	134
Table 27 - CFP control message fields	135
Table 28 - CFP control packet types	135
Table 29 - ALV control message fields.....	136
Table 30 - Keep-Alive control packet types	138
Table 31 - MUL control message fields	140
Table 32 - MUL control message types	141
Table 33 - SEC control message fields	142
Table 34 - SEC control packet types	143
Table 35 - Promotion Need MAC PDU fields.....	144
Table 36 - Beacon PDU fields.....	145
Table 37 - List of MAC primitives.....	149
Table 38 - Values of the <i>Answer</i> parameter in MAC_ESTABLISH.response primitive	152
Table 39 - Values of the <i>Result</i> parameter in MAC_ESTABLISH.confirm primitive	153
Table 40 - Values of the <i>Reason</i> parameter in MAC_RELEASE.indication primitive	154
Table 41 - Values of the <i>Answer</i> parameter in MAC_RELEASE.response primitive	154
Table 42 - Values of the <i>Result</i> parameter in MAC_RELEASE.confirm primitive	155
Table 43 - Values of the <i>Result</i> parameter in MAC_JOIN.confirm primitive.....	156
Table 44 - Values of the <i>Answer</i> parameter in MAC_ESTABLISH.response primitive.....	158
Table 45 - Values of the <i>Result</i> parameter in MAC_LEAVE.confirm primitive	159
Table 46 - Values of the <i>Result</i> parameter in MAC_DATA.confirm primitive	161
Table 47 - Values of the <i>Result</i> parameter in MLME_REGISTER.confirm primitive	163
Table 48 - Values of the <i>Result</i> parameter in MLME_UNREGISTER.confirm primitive	164
Table 49 - Values of the BCN_MODE parameter in MLME_PROMOTE.request primitive.....	165

Table 50 - Values of the Result parameter in MLME_PROMOTE.confirm primitive	166
Table 51 - Values of the <i>Result</i> parameter in MLME_DEMOTE.confirm primitive	167
Table 52 - Values of the <i>Result</i> parameter in MLME_RESET.confirm primitive	168
Table 53 - Values of the <i>status</i> parameter in MLME_GET.confirm primitive	169
Table 54 - Values of the <i>status</i> parameter in MLME_LIST_GET.confirm primitive	169
Table 55 - Values of the <i>Result</i> parameter in MLME_SET.confirm primitive	170
Table 56 - ARQ fields	199
Table 57 - Time Reference subheader fields	202
Table 58 - SAR header fields	207
Table 59 - SAR.CRC	208
Table 60 - SAR Constants	208
Table 61 - Primitive mapping between the Null SSCS primitives and the MAC layer primitives	209
Table 62 - IPV4 SSCS Table Entry	213
Table 63 - Mapping IPv4 Precedence to PRIME MAC priority	215
Table 64 - AR_REGISTER_S message format	216
Table 65 - AR_REGISTER_B message format	216
Table 66 - AR_UNREGISTER_S message format	217
Table 67 - AR_UNREGISTER_B message format	217
Table 68 - AR_LOOKUP_S message format	218
Table 69 - AR_LOOKUP_B message format	218
Table 70 - AR_MCAST_REG_S message format	219
Table 71 - AR_MCAST_REG_B message format	219
Table 72 - AR_MCAST_UNREG_S message format	219
Table 73 - AR_MCAST_UNREG_B message format	220
Table 74 - IPv4 Packet format without negotiated header compression	220
Table 75 - IPv4 Packet format with VJ header compression negotiated	221

Table 76 - Connection data sent by the initiator.....	221
Table 77 - Connection data sent by the responder.....	222
Table 78 - Connection Data sent by the Service Node.....	230
Table 79 - Connection Data sent by the Base Node.....	230
Table 80 - Node addresses table entry.....	238
Table 81 - IPv6 convergence layer table entry.....	238
Table 82 - IPv6 convergence layer multicast table entry.....	240
Table 83 - Mapping Ipv6 precedence to PRIME MAC priority.....	241
Table 84 - AR_REGISTERv6_S message format.....	242
Table 85 - AR_REGISTERv6_B message format.....	243
Table 86 - AR_UNREGISTERv6_S message format.....	243
Table 87 - AR_UNREGISTERv6_B message format.....	244
Table 88 - AR_LOOKUPv6_S message format.....	244
Table 89 - AR_LOOKUPv6_B message format.....	244
Table 90 - IPv6 Packet format without header compression.....	245
Table 91 - UDP/IPv6 Packet format with LOWPAN_IPHC1 header compression and LOWPAN_NHC.....	246
Table 92 - IPv6 Packet format with LOWPAN_IPHC negotiated header compression.....	246
Table 93 - IPv6 Unicast connection data sent by the initiator.....	247
Table 94 - IPv6 Multicast connection data sent by the initiator.....	247
Table 95 - PHY read-only variables that provide statistical information.....	255
Table 96 - PHY read-only parameters, providing information on specific implementation.....	257
Table 97 - Table of MAC read-write variables.....	258
Table 98 - Table of MAC read-only variables.....	262
Table 99 - Table of MAC read-only variables that provide functional information.....	264
Table 100 - Table of MAC read-only variables that provide statistical information.....	268
Table 101 - Table of read-only lists made available by MAC layer through management interface.....	268

Table 102 - MAC security attributes	275
Table 103 - Action PIB attributes	276
Table 104 - Applications PIB attributes	281
Table 105 - FU PIB attributes	286
Table 106 - FU State Machine	286
Table 107 - Fields of FU_INIT_REQ	292
Table 108 - Fields of FU_EXEC_REQ	294
Table 109 - Fields of FU_CONFIRM_REQ	294
Table 110 - Fields of FU_STATE_REQ	295
Table 111 - Fields of FU_KILL_REQ	295
Table 112 - Fields of FU_STATE_RSP	296
Table 113 - Fields of Exception code	297
Table 114 - Fields of FU_DATA	298
Table 115 - Fields of FU_MISS_REQ	299
Table 116 - Fields of FU_MISS_BITMAP	299
Table 117 - Fields of FU_MISS_LIST	300
Table 118 - Fields of FU_INFO_REQ	301
Table 119 - Infold possible values	302
Table 120 - Fields of FU_INFO_RSP	302
Table 121 - Fields of each entry of InfoData in FU_INFO_RSP	303
Table 122 - Management data frame fields	307
Table 123 - GET PIB Attribute request fields	308
Table 124 - GET PIB Attribute response fields	309
Table 125 - PIB List query format	311
Table 126 - PIB list response format	312
Table 127 - PHY PIB common mandatory attributes	315

Table 128 - MAC PIB comon mandatory attributes.....	315
Table 129 - Applications PIB common mandotory attributes	316
Table 130 - MAC PIB Base Node mandatory attributes	316
Table 131 - MAC PIB Service Node mandatory attributes.....	317
Table 132 - APP PIB Service Node mandatory attributes.....	318
Table 133 - Examples of CRC-8 calculated for various ASCII strings.....	319
Table 134 - Example of CRC-32.....	319
Table 135 - Header interleaving matrix.....	321
Table 136 - DBPSK(FEC ON) interleaving matrix.....	321
Table 137 - DQPSK(FEC ON) interleaving matrix.	321
Table 138 - D8PSK(FEC ON) interleaving matrix.....	322
Table 139 - Table of MAC constants.....	324
Table 140 - TYPE value assignments.....	326
Table 141 - LCID value assignments	326
Table 142 - Result values for Convergence layer primitives	326
Table 143 - Channel 1: List of frequencies used.....	330
Table 144 - Channel 2: List of frequencies used.....	331
Table 145 - Channel 3: List of frequencies used.....	333
Table 146 - Channel 4: List of frequencies used.....	334
Table 147 - Channel 5: List of frequencies used.....	336
Table 148 - Channel 6: List of frequencies used.....	338
Table 149 - Channel 7: List of frequencies used.....	339
Table 150 - Channel 8: List of frequencies used.....	341
Table 151 - Compatibility packet header fields.....	352
Table 152 - CREG control packet fields.....	355
Table 153 - CREG control packet types.....	357

Table 154 - CPRO control packet fields	359
Table 155 - CPRO control packet types	361
Table 156 - CBSI control packet fields	362
Table 157 - CBSI control message types	363
Table 158 - CFRA control packet fields.....	363
Table 159 - CFRA control packet types.....	364
Table 160 - CALV control message fields.....	364
Table 161 - Keep-Alive control packet types	364
Table 162 - Beacon PDU fields.....	365
Table 163 - PHY frame types and Payload transmission schemes	375

1 Introduction

2 This document is the technical specification for the PRIME technology.

3 1.1 Scope

4 This document specifies a PHY layer, a MAC layer and a Convergence layer for complexity-effective,
5 narrowband data transmission over electrical power lines that could be part of a Smart Grid system.

6 1.2 Overview

7 The purpose of this document is to specify a narrowband data transmission system based on OFDM
8 modulations scheme for providing mainly core utility services.

9 The specification currently describes the following:

- 10 • A PHY layer capable of achieving rates of uncoded 1Mbps (see chapter 3).
- 11 • A MAC layer for the power line environment (see chapter 4).
- 12 • A Convergence layer for adapting several specific services (see chapter 5).
- 13 • A Management Plane (see chapter 6)

14 The specification is written from the transmitter perspective to ensure interoperability between devices and
15 allow different implementations.

16 1.3 Normative references

17 The following publications contain provisions which, through reference in this text, constitute provisions of
18 this specification. At the time of publication, the editions indicated were valid. All standards are subject to
19 revision, and parties to agreements based on this Specification are encouraged to investigate the possibility
20 of applying the most recent editions of the following standards:

#	Ref.	Title
[1]	EN 50065-1:2001+A1:2010	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.
[2]	EN IEC 50065-7 Ed. 2001	Signalling on low-voltage electrical installations in the frequency range 3 kHz to 148,5 kHz. Part7: Equipment impedance.

#	Ref.	Title
[3]	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.
[4]	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).
[5]	IEC 61334-4-511 Ed. 2000	Distribution automation using distribution line carrier systems – Part 4-511: Data communication protocols – Systems management – CIASE protocol.
[6]	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.
[7]	prEN/TS 52056-8-4	Electricity metering data exchange - The DLMS/COSEM suite - Part 8-4: The PLC Orthogonal Frequency Division Multiplexing (OFDM) Type 1 profile.
[8]	IEEE Std 802-2001	IEEE Standard for Local and Metropolitan Area Networks. Overview and Architecture.
[9]	IETF RFC 768	User Datagram Protocol (UDP) [online]. Edited by J. Postel. August 1980. Available from: https://www.ietf.org/rfc/rfc768.txt
[10]	IETF RFC 791	Internet Protocol (IP) [online]. Edited by Information Sciences Institute, University of Southern California. September 1981. Available from: https://www.ietf.org/rfc/rfc791.txt
[11]	IETF RFC 793	Transmission Control Protocol (TCP) [online]. Edited by Information Sciences Institute, University of Southern California. September 1981. Available from: https://www.ietf.org/rfc/rfc793.txt
[12]	IETF RFC 1144	Compressing TCP/IP Headers for Low-Speed Serial Links [online]. Edited by V. Jacobson. February 1990. Available from: https://www.ietf.org/rfc/rfc1144.txt .

#	Ref.	Title
[13]	IETF RFC 2131	Dynamic Host Configuration Protocol (DHCP) [online]. Edited by R. Droms. March 1997. Available from: https://www.ietf.org/rfc/rfc2131.txt
[14]	IETF RFC 2460	Internet Protocol, Version 6 (IPv6) Specification [online]. Edited by S. Deering, R. Hinden. December 1998. Available from: https://www.ietf.org/rfc/rfc2460.txt
[15]	IETF RFC 3022	Traditional IP Network Address Translator (Traditional NAT) [online]. Edited by P. Srisuresh, Jasmine Networks, K. Egevang. January 2001. Available from: https://www.ietf.org/rfc/rfc3022.txt
[16]	NIST FIPS-197	Specification for the ADVANCED ENCRYPTION STANDARD (AES), http://www.csrc.nist.gov/publications/fips/fips197/fips-197.pdf
[17]	NIST SP 800-57	Recommendation for Key Management. Part 1: General (Revised). Available from http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57-Part1-revised2_Mar08-2007.pdf
[18]	NIST SP800-38A, Ed. 2001	Recommendation for Block Cipher Modes of Operation. Methods and Techniques. Available from http://csrc.nist.gov/publications/nistpubs/800-38a/sp800-38a.pdf .
[19]	IETF RFC 4191	IP version 6 addressing architecture. Available from http://tools.ietf.org/html/rfc4291 .
[20]	IETF RFC 6282	IPv6 Datagrams on IEEE 802.15.4. Available from http://tools.ietf.org/html/rfc6282 .
[21]	IETF RFC 4862	Stateless Address Configuration. Available from http://www.ietf.org/rfc/rfc4862.txt .
[22]	IETF RFC 2464	Transmission of IPv6 Packets over Ethernet Networks. Available from http://www.ietf.org/rfc/rfc4862.txt

#	Ref.	Title
[23]	NIST SP 800-108	Recommendation for Key Derivation Using Pseudorandom Functions. Available from http://csrc.nist.gov/publications/nistpubs/800-108/sp800-108.pdf
[24]	NIST SP 800-38 B	Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication. Available from http://csrc.nist.gov/publications/nistpubs/800-38B/SP_800-38B.pdf
[25]	NIST SP 800-38 C	Recommendation for Block Cipher Modes of Operation: the CCM Mode for Authentication and Confidentiality. Available from http://csrc.nist.gov/publications/nistpubs/800-38C/SP800-38C_updated-July20_2007.pdf
[26]	NIST SP 800-38 F	Recommendation for Block Cipher Modes of Operation: Methods for Key Wrapping. Available from http://dx.doi.org/10.6028/NIST.SP.800-38F
[27]	DRAFT NIST SP 800-90 C	Recommendation for Random Bit Generator (RBG) Constructions. Available from: http://csrc.nist.gov/publications/drafts/800-90/draft-sp800-90c.pdf

21 1.4 Document conventions

22 This document is divided into chapters and annexes. The document body (all chapters) is normative (except
23 for italics). The annexes may be normative or Informative as indicated for each annex.

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

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

27 Optional requirements are indicated with 'may' in the main body of this document. If an option is
28 incorporated in an implementation, it shall be applied as specified in this document.

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

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

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

1.5 Definitions

Term	Description
Band	Set of channels that may or may not be adjacent but defined for concurrent use according to channel access rules laid down in this specification.
Band Plan	Set of bands that a device is configured to operate on.
Base Node	Master Node which controls and manages the resources of a Subnetwork.
Beacon Slot	Location of the beacon PDU within a frame.
Channel	46.875KHz spectrum that may either correspond to PRIME version 1.3.6 spectrum location or any of the new extension bands defined in this version of specification.
Compliance Mode	A working mode of MAC protocol that supports existence of legacy 1.3.6 devices in a Subnetwork together with devices implementing this version of specification.
Destination Node	A Node that receives a frame.
Downlink	Data travelling in direction from Base Node towards Service Nodes
Hearing Domain	Area in which transmit signal from a device is received with some fidelity, without the need of intermediate amplification/repeating devices.
Level(PHY layer)	When used in physical layer (PHY) context, it implies the transmit power level.
Level (MAC layer)	When used in medium access control (MAC) context, it implies the position of the reference device in Switching hierarchy.
MAC frame	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 delimiter for the MAC frame.
Neighbour Node	Node A is Neighbour Node of Node B if A can directly transmit to and receive from B.

Term	Description
Node	Any one element of a Subnetwork which is able to transmit to and receive from other Subnetwork elements.
PHY frame	The set of OFDM symbols and Preamble which constitute a single PPDU
Peer	Two devices within the hearing domain of each other and having possibility or maintaining data-connectivity with each other without need of intermediate repeater / switch devices.
Preamble	The initial part of a PHY frame, used for synchronizations purposes
Registration	Process by which a Service Node is accepted as member of Subnetwork and allocated a LNID.
Service Node	Any one Node of a Subnetwork which is not a Base Node.
Source Node	A Node that sends a frame.
Subnetwork	A set of elements that can communicate by complying with this specification and share a single Base Node.
Subnetwork address	Property that universally identifies a Subnetwork. It is its Base Node EUI-48 address.
Switching	Providing connectivity between Nodes that are not Neighbour Nodes.
Unregistration	Process by which a Service Node leaves a Subnetwork.
Uplink	Data travelling in direction from Service Node towards Base Node

1.6 Abbreviations and Acronyms

Term	Description
AC	Alternating Current

Term	Description
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
CENELEC	European Committee for Electrotechnical Standardization
CFP	Contention Free Period
CID	Connection Identifier
CL	Convergence layer
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)

Term	Description
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
GK	Generation Key
GPDU	Generic MAC PDU
HCS	Header Check Sum
IEC	International Electrotechnical Committee
IEEE	Institute of Electrical and Electronics Engineers
IFFT	Inverse Fast Fourier Transform
IGMP	Internet Group Management Protocol
IPv4	Internet Protocol, Version 4

Term	Description
kbps	kilobit per second
KDIV	Key Diversifier
LCID	Local Connection Identifier
LFSR	Linear Feedback Shift Register
LLC	Logical Link Control
LNID	Local Node Identifier
LSID	Local Switch Identifier
LV	Low Voltage
LWK	Local Working Key
MAC	Medium Access Control
MK	Master Key
MLME	MAC Layer Management Entity
MPDU	MAC Protocol Data Unit
msb	Most significant bit
lsb	Least significant bit
MSPS	Million Samples Per Second
MTU	Maximum Transmission Unit
NAT	Network Address Translation

Term	Description
NID	Node Identifier
NSK	Network Secret Key
OFDM	Orthogonal Frequency Division Multiplexing
PDU	Protocol Data Unit
PHY	Physical Layer
PIB	PLC Information Base
PLC	Powerline Communications
PLME	PHY Layer Management Entity
PNPDU	Promotion Needed PDU
PPDU	PHY Protocol Data Unit
ppm	Parts per million
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

Term	Description
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
WK	Working Key

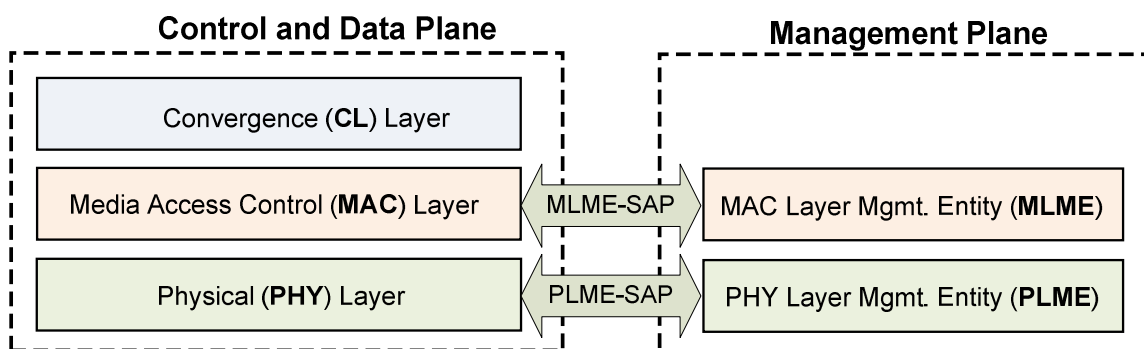
34 2 General Description

35 2.1 Introduction

36 This document is the Specification for a solution for PLC in the CENELEC A-Band using orthogonal frequency
37 division multiplexing (OFDM) modulation scheme.

38 2.2 General description of the architecture

39 Figure 1 below depicts the communication layers and the scope of this specification. This specification
40 focuses mainly on the data, control and management plane.



41

42

Figure 1 - Reference model of protocol layers used in the PRIME specification

43 The CL classifies traffic associating it with its proper MAC connection; this layer performs the mapping of any
44 kind of traffic to be properly included in MPDUs. It may also include header compression functions. Several
45 SSCs are defined to accommodate different kinds of traffic into MPDUs.

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

48 The PHY layer transmits and receives MPDUs between Neighbor Nodes using OFDM. OFDM is chosen as the
49 modulation technique because of:

- 50
- 51 • its inherent adaptability in the presence of frequency selective channels (which are common but unpredictable, due to narrowband interference or unintentional jamming);
 - 52 • its robustness to impulsive noise, resulting from the extended symbol duration and use of FEC;
 - 53 • its capacity for achieving high spectral efficiencies with simple transceiver implementations.

54 The PHY specification, described in Chapter 3, also employs a flexible coding scheme. The PHY data rates can
55 be adapted to channel and noise conditions by the MAC.

56 3 Physical layer

57 3.1 Introduction

58 This chapter specifies the PHY Entity for an OFDM modulation scheme for Power Line Communication. The
59 PHY entity uses frequencies in the band 3 kHz up to 500 kHz. The use of these frequencies is subject to
60 applicable local regulations, e.g. EN 50065 1:2001+A1:2010 in Europe or FCC part 15 in the US.

61 It is well known that frequencies below 40 kHz show several problems in typical LV power lines. For example:

- 62 • load impedance magnitude seen by transmitters is sometimes below 1Ω , especially for Base
63 Nodes located at transformers;
- 64 • colored background noise, which is always present in power lines and caused by the summation
65 of numerous noise sources with relatively low power, exponentially increases its amplitude
66 towards lower frequencies;
- 67 • meter rooms pose an additional problem, as consumer behaviors are known to have a deeper
68 impact on channel properties at low frequencies, i.e. operation of all kind of household
69 appliances leads to significant and unpredictable time-variance of both the transfer function
70 characteristics and the noise scenario.

71 Consequently, the PRIME PHY specification uses the frequency band from 41.992 kHz to 471.6796875 kHz.
72 This range is divided into eight channels, which may be used either as single independent channels or “N_{CH}”
73 of them concurrently as a unique transmission / reception band. OFDM modulation is specified in each
74 channel, with signal loaded on 97 equally spaced subcarriers, transmitted in symbols of 2240 microseconds,
75 of which 192 microseconds are comprised of a short cyclic prefix. Adjacent channels are always separated by
76 guard intervals of fifteen subcarriers (7.3 kHz). More details are provided in Annex G.

77 Differential modulations are used, with one of three possible constellations: DBPSK, DQPSK or D8PSK.

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

79 Finally, $\frac{1}{2}$ rate convolutional coding and repetition code will be used along with bit interleaving. The
80 convolutional coding, the bit interleaving and/or the repetition code can be disabled by higher layers if the
81 channel is good enough and higher throughputs are needed.

82 3.2 Overview

83 3.2.1 General


84 On the transmitter side, the PHY Layer receives a MPDU from the MAC layer and generates a PHY frame.

85 The processing of the header and the PPDU is shown in Figure 2, and consists of the following steps.

86 A CRC is appended to the PHY header (CRC for the payload is appended by the MAC layer, so no additional
87 CRC is inserted by the PHY). Next, CC is performed, if the optional FEC is enabled. The next step is scrambling,
88 which is done for both PHY header and the PPDU, irrespective of whether CC is enabled. When CC is enabled,


89 additional repetition code can be selected and, in this case, the scrambler output is repeated by a factor of
 90 four. This transmitter configuration is defined as robust mode. If CC is enabled, the scrambler output (or the
 91 repeater output in case of robust modes) is also interleaved.

92 The scrambled (and interleaved) bits are differentially modulated using a DBPSK, DQPSK or D8PSK scheme.
 93 The next step is OFDM, which comprises the IFFT block and the cyclic prefix generator. When header and
 94 data bits are input to the chain shown in Figure 2, the output of the cyclic prefix generation is a concatenation
 95 of OFDM symbols constituting the header and payload portions of the PPDU respectively. The header portion
 96 contains two or four OFDM symbols, while the payload portion contains M OFDM symbols. The value of M is
 97 signaled in the PHY header, as described in Section 3.4.3




98
99 **Figure 2 - Overview of PPDU processing**

100 Two different PHY frame formats are specified, named frame of Type A and Type B. The structure of the
 101 PRIME PHY frame of Type A is shown in Figure 3. Each PHY frame of Type A starts with a preamble lasting
 102 2.048 ms, followed by a number of OFDM symbols, each lasting 2.24 ms. The first two OFDM symbols carry
 103 the PHY frame header also referred to as the header in this specification. The header is also generated from
 104 using a process similar to the payload generation, as described in Section 3.4.3.1. The remaining M OFDM
 105 symbols carry payload, generated as described in Section 3.4.3.1. The value of M is signaled in the header
 106 and is at most equal to 63.



107
108 **Figure 3 - PHY frame of Type A format**

109 The structure of the PHY frame of Type B is shown in Figure 4. Each PHY frame of Type B starts with a
 110 preamble lasting 8.192 ms, followed by a number of OFDM symbols, each lasting 2.24 ms. The first four
 111 OFDM symbols carry the PHY frame header. The header is also generated from using a process similar to the
 112 payload generation, as described in Section 3.4.3.2. The remaining M OFDM symbols carry payload,
 113 generated as described in Section 3.4.3.2. The value of M is signaled in the header, and is at most equal to
 114 252.



115
116 **Figure 4 - PHY frame of Type B format**

117 **3.2.2 Note about backwards compatibility with PRIME v1.3.6**

118 The current version of the PRIME specification includes new features and several modifications at PHY level.
119 In order to ensure backwards compatibility with deployed PRIME devices, which shall be compliant with
120 previous PRIME specification version 1.3.6, a “PHY backwards compatibility” mechanism is described in
121 Annex J.

122 **3.3 PHY parameters**

123 Table 1 lists the frequency and timing parameters used in the PRIME PHY. These parameters are common for
124 all constellation/coding combinations.

125 **Note:** Note that throughout this document, a sampling rate of 1 MHz and 2048-point FFT sizes are defined
126 for specification convenience of the OFDM signals and are not intended to indicate a requirement on the
127 implementation

128 **Table 1 - Frequency and Timing Parameters of the PRIME PHY**

Parameter	Values	
Base Band clock (Hz)	1000000	
Subcarrier spacing (Hz)	488.28125	
Number of data subcarriers	$N_{CH} \times 84$ (header)	$N_{CH} \times 96$ (payload)
Number of pilot subcarriers	$N_{CH} \times 13$ (header)	$N_{CH} \times 1$ (payload)
FFT interval (samples)	2048	
FFT interval (μ s)	2048	
Cyclic Prefix (samples)	192	
Cyclic Prefix (μ s)	192	
Symbol interval (samples)	2240	
Symbol interval (μ s)	2240	

Parameter	Values	
Preamble period (μ s)	2048 (Type A)	8192 (Type B)

129 **Note:** $1 \leq N_{CH} \leq 8$, where “ N_{CH} ” is the number of channels as defined in Section 3.1.

130 Table 2 below shows the PHY data rate during payload transmission, and maximum MSDU length for various
 131 modulation and coding combinations. The robust modes, which include CC and repetition coding, only allow
 132 for DBPSK and DQPSK modulations. The effect of using more than one channel, as defined in Section 3.1, is
 133 represented by “ N_{CH} ” ($1 \leq N_{CH} \leq 8$).

134 **Table 2 - PHY Payload Parameters**

	DBPSK			DQPSK			D8PSK	
	On	On	Off	On	On	Off	On	Off
Convolutional Code (1/2)	On	On	Off	On	On	Off	On	Off
Repetition code	On	Off	Off	On	Off	Off	Off	Off
Information bits per subcarrier N_{BPSC}	0.5	0.5	1	1	1	2	1.5	3
Information bits per OFDM symbol N_{BPS}	$N_{CH} \times 48$	$N_{CH} \times 48$	$N_{CH} \times 96$	$N_{CH} \times 96$	$N_{CH} \times 96$	$N_{CH} \times 192$	$N_{CH} \times 144$	$N_{CH} \times 288$
Raw data rate (kbps approx)	$N_{CH} \times 5.4$	$N_{CH} \times 21.4$	$N_{CH} \times 42.9$	$N_{CH} \times 10.7$	$N_{CH} \times 42.9$	$N_{CH} \times 85.7$	$N_{CH} \times 64.3$	$N_{CH} \times 128.6$
Maximum number of payload symbols	252	63	63	252	63	63	63	63
Maximum MPDU2 length with the maximum number of payload symbols (in bits)	$N_{CH} \times 3016$	$N_{CH} \times 3016$	$N_{CH} \times 6048$	$N_{CH} \times 6040$	$N_{CH} \times 6040$	$N_{CH} \times 12096$	$N_{CH} \times 9064$	$N_{CH} \times 18144$

	DBPSK			DQPSK			D8PSK	
Maximum MPDU2 length with the maximum number of payload symbols (in bytes)	$N_{CH} \times 377$	$N_{CH} \times 377$	$N_{CH} \times 756$	$N_{CH} \times 755$	$N_{CH} \times 755$	$N_{CH} \times 1512$	$N_{CH} \times 1133$	$N_{CH} \times 2268$

135 Table 3 shows the modulation and coding scheme and the size of the header portion of the PHY frame (see
136 Section 3.4.3).

137 **Note:** The whole MPDU includes MPDU1 and MPDU2. The length of MPDU1 is defined in Section 3.4.3.1 for
138 the Type A frames and Section 3.4.3.2 for Type B frames.

139 **Table 3 - PHY Header Parameters**

	DBPSK with Header Type A	DBPSK with Header Type B
Convolutional Code (1/2)	On	On
Repetition code	Off	On
Information bits per subcarrier N_{BPSC}	0.5	0.5
Information bits per OFDM symbol N_{BPS}	$N_{CH} \times 42$	$N_{CH} \times 42$

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

142 3.4 Preamble, header and payload structure

143 3.4.1 Preamble

144 3.4.1.1 PRIME preamble Type A

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

150 The preamble of Type A, named $S(t)$, is composed by N_{CH} sub-symbols where N_{CH} is the number of channels
151 concurrently used. The set of the active channels indices is defined Ω and its i^{th} element is ω_i .

$$152 \Omega = \{ \omega \in [1, 2, \dots, 8] : \omega \text{ is an active channel} \} = \{ \omega_1, \omega_2, \dots, \omega_{N_{CH}} \}$$

153 The preamble sub-symbol $S_{SS}^c(t)$ contains a chirp signal ranging on the frequencies of channel c as defined
 154 in Annex G:

155
$$S_{SS}^c(t) = B \cdot \text{window}(t/T') \cdot \cos[2\pi(f_0^c t + 1/2 \mu_c t^2)] \quad 0 \leq t < T'$$

156 where T' is the duration of the chirp, $\mu_c = (f_f^c - f_0^c)/T'$, f_0^c and f_f^c are the start and final frequencies of
 157 the channel c , respectively. The function $\text{window}(t/T')$ is a shaping window of length T' composed by a raising
 158 roll-off region with length ro μ s a flat region (of unitary amplitude) and a decreasing roll-off region with length
 159 ro μ s. The definition of the roll-off region shape is left to individual implementations and should aim at
 160 reducing the out-of-band spectral emissions.

161 The choice of the parameter B determines the average preamble power that must be 4 dB higher than the
 162 average power of the header and payload OFDM symbols.

163 The duration T' of the sub-symbols, in μ s, is defined as follows:

164
$$T' = \frac{2048 - ro}{N_{CH}} + ro$$

165 The preamble $S(t)$ is the concatenation of the sub-symbols $S_{SS}^c(t)$ with their head and tail roll-off regions
 166 overlapped:

167
$$S(t) = \sum_{i=0}^{N_{CH}-1} S_{SS}^{a_i}(t - i \cdot (T' - ro)) \quad 0 \leq t < (T' - ro) \cdot N_{CH} + ro$$


168 To avoid rounding issues in the definition of T' , the length of the roll-off region depends on the number of
 169 active channels N_{CH} and its values are listed in Table 4.

170 **Table 4 - Roll-off region length for all N_{CH} values**

N_{CH}	ro [μ s]	N_{CH}	ro [μ s]
1	0	5	63
2	64	6	62
3	62	7	67
4	64	8	64

171 Note that when a single channel is used the roll-off regions are not present, $T' = 2048$ μ s and $S(t) \equiv S_{SS}^c(t)$
 172 .


173 Figure 5 is an example of the structure of the preamble $S(t)$ when three channels are used (channel 1,
 174 channel 3 and channel 6). In this case, $N_{CH} = 3$, $ro = 62$ μ s and $T' = 724$ μ s.



175
176 **Figure 5 - Example of the preamble structure when three channels are used**

177 **3.4.1.2 PRIME preamble Type B**

178 The preamble of Type B , named $S(t)$, is the concatenation of three preamble symbols $S_{PS}(t)$ and one
179 preamble symbol with inverted sign $-S_{PS}(t)$ as shown in Figure 6.



181 **Figure 6 - Preamble Type B structure**

182 Each preamble symbol $S_{PS}(t)$ is composed by N_{CH} sub-symbols $S_{SS}(t)$ where N_{CH} is the number of channels
183 concurrently used. The set of the active channels indices is defined Ω and its i^{th} element is ω_i .

184
$$\Omega = \{\omega \in [1,2,\dots,8]: \omega \text{ is an active channel}\} = \{\omega_1, \omega_2, \dots, \omega_{N_{CH}}\}$$

185 The sub-symbol $S_{SS}^c(t)$ contains a chirp signal ranging on the frequencies of channel c as defined in
186 Annex G:

187
$$S_{SS}^c(t) = B \cdot \text{window}(t/T') \cdot \cos[2\pi(f_0^c t + 1/2 \mu_c t^2)] \quad 0 \leq t < T'$$

188 where T' is the duration of the chirp, $\mu_c = (f_f^c - f_0^c)/T'$, f_0^c and f_f^c are the final and start frequencies of
189 the channel c , respectively. The function $\text{window}(t/T')$ is a shaping window of length T' composed by a raising
190 roll-off region with length $ro \mu\text{s}$ a flat region (of unitary amplitude) and a decreasing roll-off region with length
191 $ro \mu\text{s}$. The definition of the roll-off region shape is left to individual implementations and should aim at
192 reducing the out-of-band spectral emissions.

193 The choice of the parameter B determines the average preamble power that must be 4 dB higher than the
194 average power of the header and payload OFDM symbols.

195 The duration T' of the sub-symbols, in μs , is defined as follows:

196
$$T' = \frac{2048 - ro}{N_{CH}} + ro$$

197 The preamble symbol $S_{PS}(t)$ is the concatenation of the sub-symbols $S_{SS}^c(t)$ with their head and tail roll-
 198 off regions overlapped:

199
$$S_{PS}(t) = \sum_{i=0}^{N_{CH}-1} S_{SS}^{\omega_i}(t - i \cdot (T' - ro)) \quad 0 \leq t < (T' - ro) \cdot N_{CH} + ro$$


200 To avoid rounding issues in the definition of T' , the length of the roll-off region depends on the number of
 201 active channels N_{CH} and its values are listed in Table 5.

202 **Table 5 - Roll-off region length for all N_{CH} values**

N_{CH}	ro [μs]	N_{CH}	ro [μs]
1	0	5	63
2	64	6	62
3	62	7	67
4	64	8	64

203 Note that when a single channel is used the roll-off regions are not present, $T' = 2048 \mu s$ and $S_{PS}(t) \equiv S_{SS}^c(t)$

204 Figure 7 is an example of the structure of the preamble symbol $S_{PS}(t)$ when three channels are used
 205 (channel 1, channel 3 and channel 6). In this case, $N_{CH} = 3$, $ro = 62 \mu s$ and $T' = 724 \mu s$.



206


207 **Figure 7 - Example of each of the preamble symbol structure when three channels are used**

208 **3.4.2 Pilot structure**

209 The preamble is always followed by some OFDM symbols comprising the header. Each header symbol
 210 contains $13 \times N_{CH}$ pilot subcarriers, starting from the first subcarrier of each active channel and separated by
 211 7 data subcarriers. The pilots could be used to estimate the sampling start error and the sampling frequency
 212 offset.

213 For subsequent OFDM symbols, one pilot subcarrier is used on the first subcarrier of each active channel to
 214 provide a phase reference for frequency domain DPSK demodulation.

215 In Figure 8 pilot subcarrier allocation is shown for the eight active channels case where a 2048-point FFT is
 216 used. P_i^c is the i^{th} pilot subcarrier on the c^{th} channel and D_i^c is the i^{th} data subcarrier on the c^{th} channel.



236


237

Figure 10 - PHY frame of Type A, pilot and data subcarrier allocation (eight active channels case)

238

239 3.4.2.2 Pilot structure for PHY frames of Type B

240 In the case of PHY frame of Type B, the header is composed by four OFDM symbols. The pilot and the data
 241 subcarriers allocation for the eight active channels case is shown in Figure 11.



242

243

Figure 11 - PHY frame of Type B, pilot and data subcarrier allocation (eight active channels case)

244

3.4.3 Header and Payload

245

3.4.3.1 Header and payload for PHY frames of Type A

246

The header of Type A is composed of two OFDM symbols, which are always sent using DBPSK modulation and CC "On" (note that the repetition coding is not available for PRIME v1.3.6 devices). The payload is DBPSK, DQPSK or D8PSK modulated, depending on the configuration chosen by the MAC layer. The MAC layer may select the best modulation scheme 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

247

248

249


250

251 the best compromise between throughput and efficiency in the communication. This includes deciding
 252 whether or not CC is used.

253 **Note:** The optimization metric and the target error rate for the selection of modulation and FEC scheme is
 254 left to individual implementations

255 The first two OFDM symbols in the PPDU (corresponding to the header) are composed of $84 \times N_{CH}$ data
 256 subcarriers and $13 \times N_{CH}$ pilot subcarriers. After the header, each OFDM symbol in the payload carries $96 \times N_{CH}$
 257 data subcarriers and one pilot subcarrier. Each data subcarrier carries 1, 2 or 3 bits.

258 The bit stream from each field must be sent msb first.



259
 260 **Figure 12 - PRIME PPDU of Type A: header and payload (bits transmitted before encoding)**

- 261 • **HEADER:** The header for PRIME PPDUs of Type A comprises two OFDM symbols, containing both PHY
 262 and MAC header information. To avoid ambiguity, the MAC header is always referred to as such. The
 263 PHY header may also be referred to as just “header”. It is composed of the following fields:
 264
 - **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_C	DQPSK_C	D8PSK_C	RES	RES	RES	RES	RES	RS	RES	RES	RES

265 Where RES means “Reserved” and the suffix “_C” means CC is “On”.

- 266
 - **LEN:** defines the length of the payload (after coding) in OFDM symbols. Added by the PHY layer.
 267 If LEN is equal to 0, then the PAYLOAD symbols are not present. In this case, PAD_LEN refers to
 268 the padding bytes appended to MPDU bytes to fill the MPDU1 field.
 - 269 ○ **PAD_LEN:** defines the length of the PAD field (before coding) in bytes. The length in bits of this
 270 field depends on the number of active channels: 6 for $N_{CH} = 1$ and 9 for $N_{CH} \geq 2$. Added by the PHY
 271 layer.
 - 272 ○ **RESERVED:** contains the reserved bits for future use. The length in bits of this field depends on
 273 the number of active channels: 0 for $N_{CH} = 1$ and 5 for $N_{CH} \geq 2$.
 - 274 ○ **MPDU1:** First part of the MPDU. The length in bits of this field (MPDU1Len) depends on the
 275 number of active channels:
 276

277
$$MPDU1Len = \left\lfloor \frac{(N_{CH} \cdot 84 - 30) + 2}{8} \right\rfloor \cdot 8 - 2$$

278 The result of the above formula is reduced by 8 (1 Byte) for $N_{CH} \geq 2$.

- 279
 - where $\lfloor x \rfloor$ denotes the nearest integer towards minus infinity of x .
 - 280 ○ **CRC_Ctrl:** the $CRC_Ctrl(m)$, $m = 0..7$, contains the CRC checksum over PROTOCOL, LEN, PAD_LEN,
 281 RESERVED and MPDU1 field (PD_Ctrl). The polynomial form of PD_Ctrl is expressed as follows:

282
$$\sum_{m=0}^{69} PD_{Ctrl}(m)x^m$$

283 The checksum is calculated as follows: the remainder of the division of PD_Ctrl by the polynomial
 284 $x^8 + x^2 + x + 1$ forms CRC_Ctrl(m), where CRC_Ctrl(0) is the lsb. The generator polynomial is the
 285 well-known CRC-8-ATM. Some examples are shown in Annex A. Added by the PHY layer.

- 286 ○ FLUSHING_H: flushing bits needed for convolutional decoding. All bits in this field are set to zero
 287 to reset the convolutional encoder. Added by the PHY layer.
- 288 ○ PAD_H: Padding field. In order to ensure that the number of (coded) bits generated in the header
 289 fills an integer number of OFDM symbols, pad bits may be added to the header before encoding.
 290 All pad bits shall be set to zero. The length in bits of the PAD_H field depends on the number of
 291 active channels.

292 Table 6 resumes the length in bits of MPDU1 and PAD_H fields for different numbers of active
 293 channels.

294 **Table 6 - Length in bits of MPDU1 and PAD_H fields in the PHY frame header of Type A for all possible values of N_{CH}**

N _{CH}	MPDU1	PAD_H
1	54	0
2	126	4
3	214	0
4	294	4
5	382	0
6	462	4
7	550	0
8	630	4

- 295 ● PAYLOAD:
 - 296 ○ MPDU2: Second part of the MPDU.
 - 297 ○ FLUSHING_P: flushing bits needed for convolutional decoding. All bits in this field are set to zero
 298 to reset the convolutional encoder. This field only exists when CC is “On”.
 - 299 ○ PAD: Padding field. In order to ensure that the number of (coded) bits generated in the payload
 300 fills an integer number of OFDM symbols, pad bits may be added to the payload before encoding.
 301 All pad bits shall be set to zero.

302 The MPDU is included in the MPDU1 and MPDU2 fields using the following logic. The first 2 bits of the MPDU
 303 are discarded for alignment purposes. The next 54 bits of the MPDU are included in the MPDU1 field. The
 304 remaining bits of the MPDU are included in the MPDU2 field. It is a work of higher layers not to use the first
 305 two bits of the MPDU as they will not be transmitted or received by the PHY layer. In reception these first
 306 non-transmitted bits will be considered as 0.


307 **3.4.3.2 Header and payload for PHY frames of Type B**

308 The header is composed of four OFDM symbols, which are always sent using DBPSK modulation, CC “On” and
 309 repetition coding “On”. However the payload is DBPSK, DQPSK or D8PSK modulated, depending on the
 310 configuration by the MAC layer. The MAC layer may select the best modulation scheme using information
 311 from errors in previous transmissions to the same receiver(s), or by using the SNR feedback. Thus, the system
 312 will then configure itself dynamically to provide the best compromise between throughput and efficiency in
 313 the communication. This includes deciding whether or not CC and repetition coding are used.

314 **Note:** *The optimization metric and the target error rate for the selection of modulation and FEC scheme is*
 315 *left to individual implementations*

316 The first four OFDM symbols in the PPDU (corresponding to the header) are composed of 84×N_{CH} data
 317 subcarriers and 13×N_{CH} pilot subcarriers. After the header, each OFDM symbol in the payload carries 96×N_{CH}
 318 data subcarriers and N_{CH} pilot subcarriers. Each data subcarrier carries 1, 2 or 3 bits.

319 The bit stream from each field must be sent msb first.



320
321 **Figure 13 - PRIME PPDU of Type B: header and payload (bits transmitted before encoding)**

- 322 • HEADER: The PHY header is composed of the following fields:
- 323 ○ 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_C	DQPSK_C	D8PSK_C	RES	RES	RES	RES	RES	R_DBPSK	R_DQPSK	RES	RES

324 Where RES means “Reserved”, the suffix “_C” means CC is “On” and the prefix R_ means CC and
 325 repetition are “On”.
 326

- 327 ○ LEN: defines the length of the payload (after coding) in OFDM symbols. Added by the PHY layer.
- 328 ○ PAD_LEN: defines the length of the PAD field (before coding) in bytes. If LEN is equal to 0 the
 329 PAYLOAD symbols are not present, in this case PAD_LEN refers to the padding bytes appended to
 330 MPDU bytes to fill the MPDU1 field. Added by the PHY layer.
- 331 ○ RESERVED: contains the reserved bits for future use.
- 332 ○ CRC_Ctrl: the CRC_Ctrl(m), m = 0..11, contains the CRC checksum over PROTOCOL, LEN, PAD_LEN
 333 and RESERVED field (PD_Ctrl). The polynomial form of PD_Ctrl is expressed as follows:

$$\sum_{m=0}^{23} PD_{Ctrl}(m)x^m$$

334
 335 The checksum is calculated as follows: the remainder of the division of PD_Ctrl by the polynomial
 336 $x^{12} + x^{11} + x^3 + x^2 + x + 1$ forms CRC_Ctrl(m), where CRC_Ctrl(0) is the lsb. Some examples are
 337 shown in Annex A. Added by the PHY layer.

- 338 ○ MPDU1: First part of the MPDU. The length in bits of this field (MPDU1Len) is a multiple of 8 and
 339 it depends on the number of active channels:

$$MPDU1Len = \left\lfloor \frac{(N_{CH} - 1) \cdot 84 \cdot \frac{1}{2}}{8} \right\rfloor \cdot 8$$

where $\lfloor x \rfloor$ denotes the nearest integer towards minus infinity of x .

- 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.
- PAD_H: Padding field. In order to ensure that the number of (coded) bits generated in the header fills an integer number of OFDM symbols, pad bits may be added to the header before encoding. All pad bits shall be set to zero. The length in bits of PAD_H field (PAD_HLen) depends on the number of active channels:

$$PAD_HLen = (N_{CH} - 1) \cdot 84 \cdot \frac{1}{2} - MPDU1Len$$

Table 7 resumes the length of MPDU1 and PAD_H fields for different numbers of active channels.

Table 7 - Length in bits of MPDU1 and PAD_H fields in the PHY frame header of Type B for all possible values of N_{CH}

N_{CH}	MPDU1	PAD_H
1	0	0
2	40	2
3	80	4
4	120	6
5	168	0
6	208	2
7	248	4
8	288	6


Note that on reception of a PPDU with a correct CRC_Ctrl but with PROTOCOL with reserved values, or any of the reserved bits being "1", the receiver should consider that the payload contains LEN symbols, and should be able to discard the PDU considering the channel busy.

- PAYLOAD:

- MPDU2: Second part of the MPDU.
- 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 CC is "On".
- PAD: Padding field. In order to ensure that the number of (coded) bits generated in the payload fills an integer number of OFDM symbols, pad bits may be added to the payload before encoding. All pad bits shall be set to zero.

361 **3.5 Convolutional encoder**

362 The uncoded bit stream may go through convolutional coding to form the coded bit stream. The
 363 convolutional encoder is 1/2 rate with constraint length K = 7 and code generator polynomials 1111001 and
 364 1011011. At the start of every PPDU transmission, the encoder state is set to zero. As seen in Figure 12 and
 365 Figure 13, six zeros are inserted at the end of the header information bits to flush the encoder and return the
 366 state to zero. Similarly, if convolutional encoding is used for the payload, eight zeros bits are again inserted
 367 at the end of the input bit stream to ensure the encoder state returns to zero at the end of the payload. The
 368 block diagram of the encoder is shown in Figure 14.



369
 370 **Figure 14 - Convolutional encoder**


371 **3.6 Scrambler**

372 The scrambler block randomizes the bit stream, so it reduces the crest factor at the output of the IFFT when
 373 a long stream of zeros or ones occurs in the header or payload bits after coding (if any). Scrambling is always
 374 performed regardless of the modulation and coding configuration.

375 The scrambler block performs a xor of the input bit stream by a pseudo noise sequence pn, obtained by cyclic
 376 extension of the 127-element sequence given by:

377 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,
 378 0,1,1,0,0,0,0,0,1,1,0,0,1,1,0,1,0,1,0,0,1,1,1,0,0,1,1,1,1,0,1,1,0,1,0,0,0,0,1,0,1,0,1,0,1,1,1,1,1,0,1,0,0,1,0,1,0,0,
 379 ,0,1,1,0,1,1,1,0,0,0,1,1,1,1,1,1}

380 **Note:** The above 127-bit sequence can be generated by the LFSR defined in Figure 9 when the “all ones” initial
 381 state is used.



382

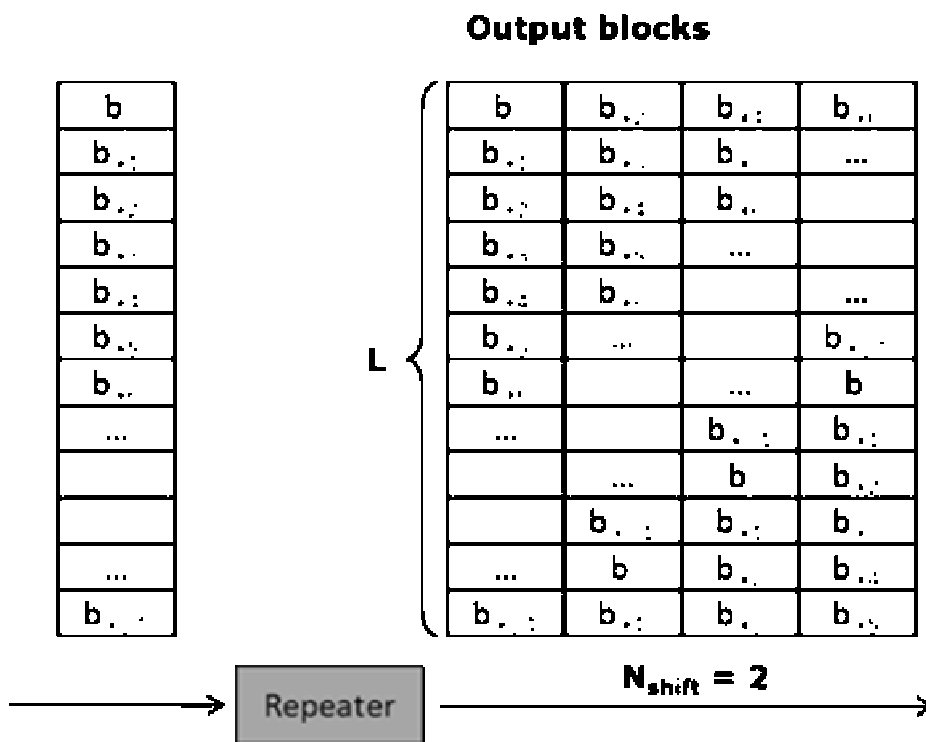
383

Figure 15 - LFSR for use in the scrambler block

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

385 3.7 Repeater

386 The repeater block introduces both time diversity and frequency diversity to the transmitted bits repeating
 387 a bit sequence four times with the aim of increasing the communication robustness. The repeater is enabled
 388 only when robust modes are used. Figure 16 shows the behavior of the repeater.



389

390

Figure 16 - Example of repeater block using a shift value = 2

391 The transmitted bit sequence b_i, b_{i+1}, \dots is divided into blocks of length L corresponding to the number of bits
 392 transmitted into one OFDM symbol according to the used transmission mode. L is equal to $84 \times N_{CH}$ for the
 393 header, $96 \times N_{CH}$ for the payload using robust DBPSK and $192 \times N_{CH}$ for the payload using robust DQPSK, where
 394 N_{CH} is the number of channels concurrently used. Each block of L bits is repeated four times at the repeater
 395 output. Furthermore, the bits of each replicated block are obtained introducing a cyclic shift of N_{shift} to the
 396 bits of the previous block (the first output block always corresponds to the input block). N_{shift} depends on the
 397 transmission mode and its values are listed in Table 8.

Table 8 - Shift values for the Robust modes

Transmission mode	N_{shift}
Robust DBPSK (header)	2
Robust DBPSK (payload)	2
Robust DQPSK (payload)	4

399 3.8 Interleaver

400 Because of the frequency fading (narrowband interference) of typical power line channels, OFDM subcarriers
 401 are generally received at different amplitudes. Deep fades in the spectrum may cause groups of subcarriers
 402 to be less reliable than others, thereby causing bit errors to occur in bursts rather than be randomly scattered.
 403 If (and only if) coding is used as described in 3.4.3, interleaving is applied to randomize the occurrence of bit
 404 errors prior to decoding. At the transmitter, the coded bits are permuted in a certain way, which makes sure
 405 that adjacent bits are separated by several bits after interleaving.

406 Let $N_{\text{CBPS}} = 2 \times N_{\text{BPS}}$ be the number of coded bits per OFDM symbol in the cases convolutional coding is used.
 407 All coded bits must be interleaved by a block interleaver with a block size corresponding to N_{CBPS} . The
 408 interleaver ensures that adjacent coded bits are mapped onto non-adjacent data subcarriers. Let $v(k)$, with k
 409 = $0, 1, \dots, N_{\text{CBPS}} - 1$, be the coded bits vector at the interleaver input. $v(k)$ is transformed into an interleaved
 410 vector $w(i)$, with $i = 0, 1, \dots, N_{\text{CBPS}} - 1$, by the block interleaver as follows:

$$411 \quad w((N_{\text{CBPS}}/s) \times (k \bmod s) + \text{floor}(k/s)) = v(k) \quad k = 0, 1, \dots, N_{\text{CBPS}} - 1$$

412 The value of s is determined by the number of coded bits per subcarrier, $N_{\text{CBPSC}} = 2 \times N_{\text{BPSC}}$. N_{CBPSC} is related to
 413 N_{CBPS} such that $N_{\text{CBPS}} = 96 \times N_{\text{CBPSC}} \times N_{\text{CH}}$ (payload) and $N_{\text{CBPS}} = 84 \times N_{\text{CBPSC}} \times N_{\text{CH}}$ (header), where N_{CH} is the number
 414 of channels concurrently used.

$$415 \quad s = 8 \times (1 + \text{floor}(N_{\text{CBPSC}}/2)) \text{ for the payload and}$$

$$416 \quad s = 7 \text{ for the header.}$$

417 At the receiver, the de-interleaver performs the inverse operation. Hence, if $w'(i)$, with $i = 0, 1, \dots, N_{\text{CBPS}} - 1$, is
 418 the de-interleaver vector input, the vector $w'(i)$ is transformed into a de-interleaved vector $v'(k)$, with $k =$
 419 $0, 1, \dots, N_{\text{CBPS}} - 1$, by the block de-interleaver as follows:

$$420 \quad v'(s \times i - (N_{\text{CBPS}} - 1) \times \text{floor}(s \times i / N_{\text{CBPS}})) = w'(i) \quad i = 0, 1, \dots, N_{\text{CBPS}} - 1$$

421 Descriptive tables showing index permutations can be found in Annex C for reference.


422 Note that the interleaver parameters k and N_{CBPS} do not depend on the presence of the repetition encoding
 423 and their values remain the same for coded DBPSK (or coded DQPSK) and robust DBPSK (or robust DQPSK).

424 3.9 Modulation

425 The PDU payload is modulated as a multicarrier differential phase shift keying signal with one pilot
 426 subcarrier and $96 \times N_{CH}$ data subcarriers that comprise $96 \times N_{CH}$, $192 \times N_{CH}$ or $288 \times N_{CH}$ bits per symbol. The
 427 header is modulated DBPSK with $13 \times N_{CH}$ pilot subcarriers and $84 \times N_{CH}$ data subcarriers that comprise $84 \times N_{CH}$
 428 bits per symbol.

429 The bit stream coming from the interleaver is divided into groups of B bits where the first bit of the group of
 430 B is the most significant bit (msb).

431 First of all, frequency domain differential modulation is performed. Figure 17 shows the DBPSK, DQPSK and
 432 D8PSK mapping:



433
 434 **Figure 17 - DBPSK, DQPSK and D8PSK mapping**

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

436
$$s_k = A e^{j\theta_k}$$

437 Where:


- 438 • k is the frequency index representing the k^{th} subcarrier in an OFDM symbol. $k = 1$ corresponds to
 439 the phase reference pilot subcarrier.
- 440 • s_k is the modulator output (a complex number) for the k^{th} given subcarrier.
- 441 • θ_k stands for the absolute phase of the modulated signal, and is obtained as follows:
- 442 • $\theta_k = (\theta_{k-1} + (2\pi / P) \Delta b_k) \text{ mod } 2\pi$
- 443 • This equation applies for $k > 1$ in the payload, the $k = 1$ subcarrier being the phase reference pilot.
 444 When the header is transmitted, the pilot allocated in the k^{th} subcarrier is used as a phase
 445 reference for the data allocated in the $(k+1)^{\text{th}}$ subcarrier.
- 446 • $\Delta b_k \in \{0, 1, \dots, P-1\}$ represents the information coded in the phase increment, as supplied by
 447 the constellation encoder.
- 448 • $P = 2, 4, \text{ or } 8$ in the case of DBPSK, DQPSK or D8PSK, respectively.

449
450
451

- A is a shaping parameter and represents the ring radius from the center of the constellation. The value of A determines the power in each subcarrier and hence the average power transmitted in the header and payload symbols.

452
453

If a complex 2048-point IFFT is used, the $96 \times N_{CH}$ subcarriers shall be mapped as shown in Figure 18. The symbol * represents complex conjugate.



454

455

Figure 18 - Subcarrier Mapping

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

457 3.10 Electrical specification of the transmitter

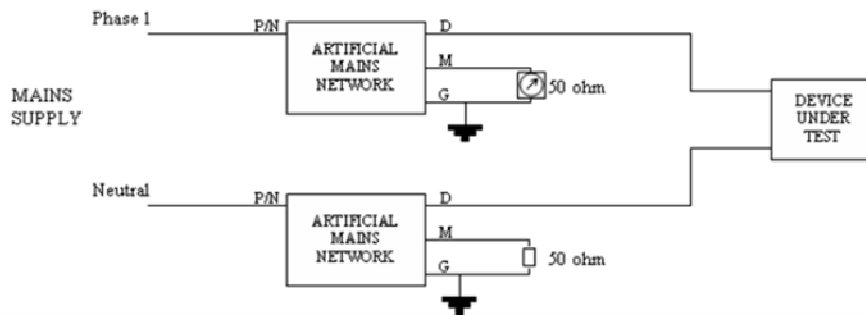
458 3.10.1 General

459 The following requirements establish the minimum technical transmitter requirements for interoperability,
460 and adequate transmitter performance.

461 3.10.2 Transmit PSD

462 Transmitter specifications will be measured according to the following conditions and set-up.

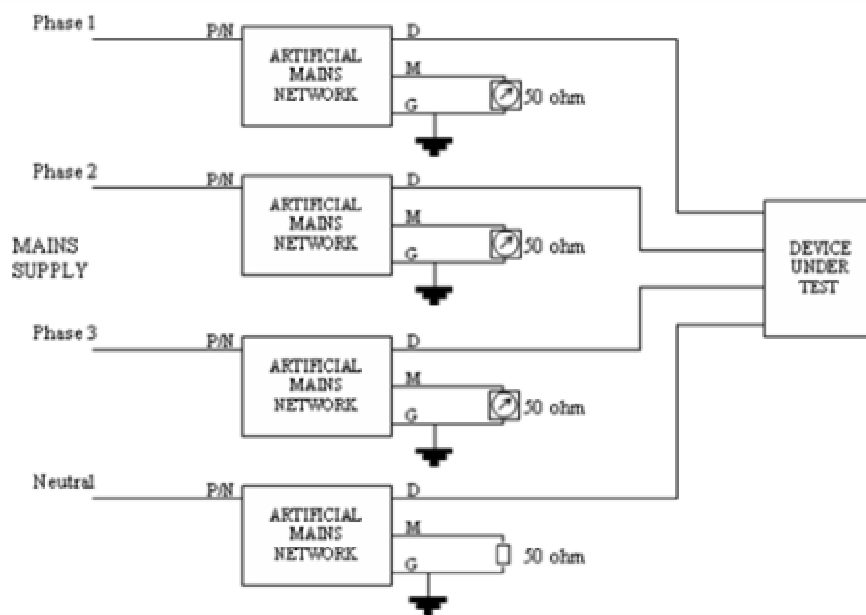
463 For single-phase devices, the measurement shall be taken on either the phase or neutral connection
464 according to Figure 19.



465

466 Figure 19 - Measurement set up (single-phase)


467 For three-phase devices which transmit on all three phases simultaneously, measurements shall be taken in
468 all three phases as per Figure 20. No measurement is required on the neutral conductor.



469

470 Figure 20 - Measurement set up (three-phase)

471 The artificial mains network in Figure 19 and Figure 20 is shown in Figure 21. It is based on EN 50065-1:2001.
 472 The 33uF capacitor and 1Ω resistor have been introduced so that the network has an impedance of 2Ω in the
 473 frequency band of interest.



474

475

Figure 21 - Artificial mains network

476 All transmitter output voltages are specified as the voltage measured at the line Terminal with respect to the
 477 neutral Terminal. Accordingly, values obtained from the measuring device must be increased by 6 dB (voltage
 478 divider of ratio $\frac{1}{2}$).

479 All devices will be tested to comply with PSD requirements over the full temperature range, which depends
 480 on the type of Node:

- 481 • Base Nodes in the range -40°C to +70°C
- 482 • Service Nodes in the range -25°C to +55°C

483 All tests shall be carried out under normal traffic load conditions.

484 In all cases, the PSD must be compliant with the regulations in force in the country where the system is used.

485 When driving only one phase, the power amplifier shall be capable of injecting a final signal level in the
 486 transmission Node (S1 parameter) of 120dBμVrms (1 Vrms). This could be in one of two scenarios: either the
 487 DUT is connected to a single phase as shown in Figure 19; or the DUT is connected to three phases as shown
 488 in Figure 20, but drives only one phase at a time. In both cases, connection is through the AMN of Figure 21.


489 For three-phase devices injecting simultaneously into all three phases, the final signal level shall be
 490 114dBμVrms (0.5Vrms).

491 **Note 1:** In all the above cases, note the measurement equipment has some insertion loss. Specifically, in the
 492 single-phase, configuration, the measured voltage is 6 dB below the injected signal level, and will equal 114
 493 dBuV when the injected signal level is 120 dBuV. Similarly, when connected to three phases, the measured
 494 signal level will be 12 dB below the injected signal level. Thus, a 114 dBuV signal injected into three phases
 495 being driven simultaneously, will be measured as 102 dBuV on any of the three meters of Figure 20.

496 **Note 2:** Regional restrictions may apply, ex., on the reactive power drawn from a meter including a PRIME
 497 modem. These regulations could affect the powerline interface, and should be accounted for.

498 **3.10.3 Error Vector Magnitude (EVM)**

499 The quality of the injected signal with regard to the artificial mains network impedance must be measured in
500 order to validate the transmitter device. Accordingly, a vector analyzer that provides EVM measurements
501 (EVM meter) shall be used, see Annex B for EVM definition. The test set-up described in Figure 19 and Figure
502 20 shall be used in the case of single-phase devices and three-phase devices transmitting simultaneously on
503 all phases, respectively.



504

505

Figure 22 - EVM meter (block diagram)

506 The EVM meter must include a Band Pass Filter with an attenuation of 40 dB at 50 Hz that ensures anti-
507 aliasing for the ADC. The minimum performance of the ADC is 1MSPS, 14-bit ENOB. The ripple and the group
508 delay of the band pass filter must be accounted for in EVM calculations.

509 **3.10.4 Conducted disturbance limits**

510 Regional regulations may apply. For instance, in Europe, transmitters shall comply with the maximum
511 emission levels and spurious emissions defined in EN50065-1:2001 for conducted emissions in AC mains in
512 the bands 3 kHz to 9 kHz and 95 kHz to 30 MHz. European regulations also require that transmitters and
513 receivers shall comply with impedance limits defined in EN50065-7:2001 in the range 3 kHz to 148.5 kHz.

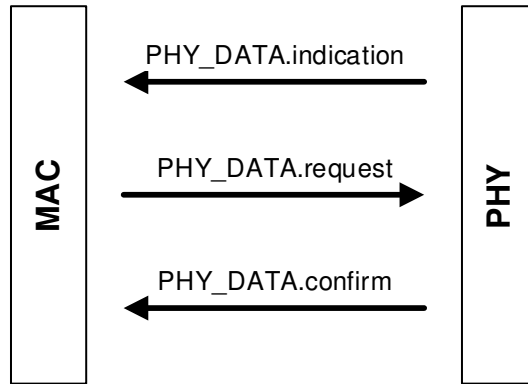
514 **3.11 PHY service specification**

515 **3.11.1 General**

516 PHY shall have a single 20-bit free-running clock incremented in steps of 10 μ s. The clock counts from 0 to
517 1048575 then overflows back to 0. As a result the period of this clock is 10.48576 seconds. The clock is never
518 stopped nor restarted. Time measured by this clock is the one to be used in some PHY primitives to indicate
519 a specific instant in time.

520 **3.11.2 PHY Data plane primitives**

521 **3.11.2.1 General**



522

523

Figure 23 - Overview of PHY primitives

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

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

528 **3.11.2.2 PHY_DATA.request**

529 **3.11.2.2.1 Function**

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

533 **3.11.2.2.2 Structure**

534 The semantics of this primitive are as follows:

535 PHY_DATA.request{MPDU, Length, Level, Type, Scheme, Scheduled, Time}.

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

539 The *Length* parameter specifies the length of MPDU in bytes. Length is 2 bytes long.

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

542 0: Maximal output level (MOL)

543 1: MOL -3 dB

544 2: MOL -6 dB

545 ...

546 7: MOL -21 dB

547 The *Type* parameter specifies the PHY frame type which should be used for the transmission: 0: PHY frame
548 Type A 1: PHY frame Type B.

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

551 0: DBPSK

552 1: DQPSK

553 2: D8PSK

554 3: Not used

555 4: DBPSK + Convolutional Code

556 5: DQPSK + Convolutional Code

557 6: D8PSK + Convolutional Code

558 7-11: Not used

559 12: Robust DBPSK

560 13: Robust DQPSK

561 14-15: Not used

562 If *Scheduled* is false, the transmissions shall start as soon as possible. If *Scheduled* is true, the *Time* parameter
563 is taken into account. The *Time* parameter specifies the instant in time in which the MPDU has to be
564 transmitted. It is expressed in 10s of μ s and may take values from 0 to $2^{20}-1$.

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

569 **3.11.2.2.3 Use**

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

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

575 3.11.2.3 PHY_DATA.confirm

576 3.11.2.3.1 Function

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

580 3.11.2.3.2 Structure

581 The semantics of this primitive are as follows:

582 PHY_DATA.confirm{*Result*}.

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

586 0: Success.

587 1: Too late. Time for transmission is past.

588 2: Invalid *Length*.

589 3: Invalid *Scheme*.

590 4: Invalid *Level*.

591 5: Buffer overrun.

592 6: Busy channel.

593 7-255: Proprietary.

594 3.11.2.3.3 Use

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

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

598 3.11.2.4 PHY_DATA.indication

599 3.11.2.4.1 Function

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

601 3.11.2.4.2 Structure

602 The semantics of this primitive are as follows:

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

604 The *PSDU* parameter specifies the PHY service data unit as received by the local PHY layer entity. It is
605 mandatory for implementations to byte-align MPDU across the PHY-SAP. For Type A frames, this implies 2
606 extra bits (due to the non-byte-aligned nature of the MAC layer Header) to be located at the beginning of the
607 header.

608 The *Length* parameter specifies the length of received PSDU in bytes. Length is 2 bytes long.

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

611 0: ≤ 70 dBuV

612 1: ≤ 72 dBuV

613 2: ≤ 74 dBuV

614 ...

615 15: > 98 dBuV

616 The *Type* parameter specifies the PHY frame type with which PSDU is received: 0: PHY frame Type A 1: PHY
617 frame Type B.

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

620 0: DBPSK

621 1: DQPSK

622 2: D8PSK

623 3: Not used

624 4: DBPSK + Convolutional Code

625 5: DQPSK + Convolutional Code

626 6: D8PSK + Convolutional Code

627 7-11: Not used

628 12: Robust DBPSK

629 13: Robust DQPSK

630 14-15: Not used

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


632 **3.11.2.4.3 Use**

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

635 **3.11.3 PHY Control plane primitives**

636 **3.11.3.1 General**

637 Figure 24 shows the generate structure of PHY control plane primitives. Each primitive may have "set", "get"
 638 and "confirm" fields. Table 4 below lists the control plane primitives and the fields associated with each of
 639 them. Each row is a control plane primitive. An "X" in a column indicates that the associated field is used in
 640 the primitive described in that row.



641

642

Figure 24 - Overview of PHY Control Plane Primitives

643

Table 9 - Fields associated with PHY Control Plane Primitives

Field	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

644 **3.11.3.2 PHY_AGC.set**

645 **3.11.3.2.1 Function**

646 The PHY_AGC.set primitive is passed to the PHY layer entity by the MAC layer entity to set the Automatic
 647 Gain Mode of the PHY layer.

648 **3.11.3.2.2 Structure**

649 The semantics of this primitive are as follows:

650 $\text{PHY_AGC.set}\{Mode, Gain\}.$

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

653 0: Auto;

654 1: Manual.

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

656 0: *min_gain* dB;

657 1: *min_gain + step* dB;

658 2: *min_gain + 2*step* dB;

659 ...

660 N-1: *min_gain + (N-1)*step* dB.

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

663 **3.11.3.2.3 Use**

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

665 **3.11.3.3 PHY_AGC.get**

666 **3.11.3.3.1 Function**

667 The PHY_AGC.get primitive is passed to the PHY layer entity by the MAC layer entity to get the Automatic
668 Gain Mode of the PHY layer.

669 **3.11.3.3.2 Structure**

670 The semantics of this primitive are as follows:

671 $\text{PHY_AGC.get}\{ \}.$

672 **3.11.3.3.3 Use**

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

675 **3.11.3.4 PHY_AGC.confirm**

676 **3.11.3.4.1 Function**

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

679 **3.11.3.4.2 Structure**

680 The semantics of this primitive are as follows:

681
$$\text{PHY_AGC.confirm } \{Mode, Gain\}.$$

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

684 0: Auto;

685 1: Manual.

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

687 0: *min_gain* dB;

688 1: *min_gain* + *step* dB;

689 2: *min_gain* + 2**step* dB;

690 ...

691 N-1: *min_gain* + (N-1)**step* dB.

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

694 **3.11.3.5 PHY_Timer.get**

695 **3.11.3.5.1 Function**

696 The PHY_Timer.get primitive is passed to the PHY layer entity by the MAC layer entity to get the current PHY
697 time.

698 **3.11.3.5.2 Structure**

699 The semantics of this primitive are as follows:

700
$$\text{PHY_Timer.get } \{ \}.$$

701 **3.11.3.5.3 Use**

702 The primitive is generated by the MAC layer to know the current PHY time.

703 3.11.3.6 PHY_Timer.confirm

704 3.11.3.6.1 Function

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

707 3.11.3.6.2 Structure

708 The semantics of this primitive are as follows:

709
$$\text{PHY_Timer.confirm } \{Time\}.$$

710 The *Time* parameter is specified in 10s of microseconds. It may take values of between 0 and $2^{20}-1$.

711 3.11.3.7 PHY_CD.get

712 3.11.3.7.1 Function

713 The PHY_CD.get primitive is passed to the PHY layer entity by the MAC layer entity to look for the carrier
714 detect signal. The carrier detection algorithm shall be based on preamble detection and header recognition
715 (see Section 3.4).

716 3.11.3.7.2 Structure

717 The semantics of this primitive are as follows:

718
$$\text{PHY_CD.get } \{ \}.$$

719 3.11.3.7.3 Use

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

722 3.11.3.8 PHY_CD.confirm

723 3.11.3.8.1 Function

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

726 3.11.3.8.2 Structure

727 The semantics of this primitive are as follows:

728
$$\text{PHY_CD.confirm } \{cd, rssi, Time, header\}.$$

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

730 0: no carrier detected;

731 1: carrier detected.

732 The *rssl* parameter is the Received Signal Strength Indication, not including the noise power. One of the RSSI
733 estimator examples is shown in Annex B, but it is implementation specific. It is only relevant when *cd* equals
734 1. It may take one of sixteen values:

735 0: ≤ 70 dBuV;

736 1: ≤ 72 dBuV;

737 2: ≤ 74 dBuV;

738 ...

739 15: > 98 dBuV.

740 The *Time* parameter indicates the instant at which the present PPDU will finish. It is only relevant when *cd*
741 equals 1. When *cd* equals 0, *Time* parameter will take a value of 0. If *cd* equals 1 but the duration of the whole
742 PPDU is still not known (i.e. the header has not yet been processed), *header* parameter will take a value of 1
743 and *time* parameter will indicate the instant at which the header will finish, specified in 10s of microseconds.
744 In any other case the value of *Time* parameter is the instant at which the present PPDU will finish, and it is
745 specified in 10s of microseconds. *Time* parameter refers to an absolute point in time so it is referred to the
746 system clock.

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

748 1: if a preamble has been detected but the duration of the whole PPDU is not yet known from
749 decoding the header;

750 0: in any other case.

751 3.11.3.9 PHY_NL.get

752 3.11.3.9.1 Function

753 The PHY_NL.get primitive is passed to the PHY layer entity by the MAC layer to get the noise floor level value.
754 One of the noise estimator examples is shown in Annex B, but it is implementation specific.

755 3.11.3.9.2 Structure

756 The semantics of this primitive are as follows:

757 PHY_NL.get {}.

758 3.11.3.9.3 Use

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

760 **3.11.3.10 PHY_NL.confirm**

761 **3.11.3.10.1 Function**

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

764 **3.11.3.10.2 Structure**

765 The semantics of this primitive are as follows:

766
$$\text{PHY_NL.confirm } \{noise\}.$$

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

768 0: ≤ 50 dBuV;

769 1: ≤ 53 dBuV;

770 2: ≤ 56 dBuV;

771 ...

772 15: > 92 dBuV.

773 **3.11.3.11 PHY_SNR.get**

774 **3.11.3.11.1 Function**

775 The PHY_SNR.get primitive is passed to the PHY layer entity by the MAC layer entity to get the value of the
776 Signal to Noise Ratio, defined as the ratio of measured received signal level to noise level of last received
777 PPDU. The calculation of the SNR is described in Annex B.

778 **3.11.3.11.2 Structure**

779 The semantics of this primitive are as follows:

780
$$\text{PHY_SNR.get } \{\}.$$

781 **3.11.3.11.3 Use**

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

784 **3.11.3.12 PHY_SNR.confirm**

785 **3.11.3.12.1 Function**

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

788 **3.11.3.12.2 Structure**

789 The semantics of this primitive are as follows:

790 $\text{PHY_SNR.confirm}\{SNR\}$.

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

794 0: ≤ 0 dB;

795 1: ≤ 3 dB;

796 2: ≤ 6 dB;

797 ...

798 7: > 18 dB.

799 **3.11.3.13 PHY_ZCT.get**

800 **3.11.3.13.1 Function**

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

803 **3.11.3.13.2 Structure**

804 The semantics of this primitive are as follows:

805 $\text{PHY_ZCT.get}\{\}$.

806 **3.11.3.13.3 Use**

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

809 **3.11.3.14 PHY_ZCT.confirm**

810 **3.11.3.14.1 Function**

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

813 **3.11.3.14.2 Structure**

814 The semantics of this primitive are as follows:

815 $\text{PHY_ZCT.confirm}\{Time\}$.

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

817 **3.11.4 PHY Management primitives**

818 **3.11.4.1 General**

819 PHY layer management primitives enable the conceptual PHY layer management entity to interface to upper
820 layer management entities. Implementation of these primitives is optional. Please refer to Figure 24 to see
821 the general structure of the PHY layer management primitives.

822 **3.11.4.2 PLME_RESET.request**

823 **3.11.4.2.1 Function**

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

829 **3.11.4.2.2 Structure**

830 The semantics of this primitive are as follows:

831 `PLME_RESET.request{}`.

832 **3.11.4.2.3 Use**

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

835 **3.11.4.3 PLME_RESET.confirm**

836 **3.11.4.3.1 Function**

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

839 **3.11.4.3.2 Structure**

840 The semantics of this primitive are as follows:

841 `PLME_RESET.confirm{Result}`.

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

843 0: Success;

844 1: Failure. The requested reset failed due to internal implementation issues.

845 **3.11.4.3.3 Use**

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

847 **3.11.4.4 PLME_SLEEP.request**

848 **3.11.4.4.1 Function**

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

852 **3.11.4.4.2 Structure**

853 The semantics of this primitive are as follows:

854 `PLME_SLEEP.request{}`.

855 **3.11.4.4.3 Use**

856 Although this specification pertains to communication over power lines, it may still be objective of some
857 applications to optimize their power consumption. This primitive is designed to help those applications
858 achieve this objective.

859 **3.11.4.5 PLME_SLEEP.confirm**

860 **3.11.4.5.1 Function**

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

863 **3.11.4.5.2 Structure**

864 The semantics of this primitive are as follows:

865 `PLME_SLEEP.confirm{Result}`.

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

867 0: Success;

868 1: Failure. The requested sleep failed due to internal implementation issues;

869 2: PHY layer is already in sleep state.

870 **3.11.4.5.3 Use**

871 The primitive is generated in response to a PLME_SLEEP.request

872 **3.11.4.6 PLME_RESUME.request**

873 **3.11.4.6.1 Function**

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

876 **3.11.4.6.2 Structure**

877 The semantics of this primitive are as follows:

878 `PLME_RESUME.request{}`.

879 **3.11.4.6.3 Use**

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

883 **3.11.4.7 PLME_RESUME.confirm**

884 **3.11.4.7.1 Function**

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

887 **3.11.4.7.2 Structure**

888 The semantics of this primitive are as follows:

889 `PLME_RESUME.confirm{Result}`.

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

891 0: Success;

892 1: Failure. The requested resume failed due to internal implementation issues;

893 2: PHY layer is already in fully functional state.

894 **3.11.4.7.3 Use**

895 The primitive is generated in response to a `PLME_RESUME.request`

896 **3.11.4.8 PLME_TESTMODE.request**

897 **3.11.4.8.1 Function**

898 The `PLME_TESTMODE.request` primitive is invoked to enter the PHY layer to a test mode (specified by the
899 mode parameter). A valid packet is transmitted and the PSDU will contain a defined reference: dummy 54-
900 bit MAC header, message "PRIME IS A WONDERFUL TECHNOLOGY" (note the blank spaces so it represents
901 240 uncoded bits in ASCII format) concatenated as many times as needed to make it 256bytes. The last eight
902 bits will be substituted for eight flushing bits set to zero. Following receipt of this primitive, the PHY layer
903 should complete any pending transmissions in its buffer before entering the requested Test mode.

904 **3.11.4.8.2 Structure**

905 The semantics of this primitive are as follows:

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

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

908 0: stop test mode and return to normal functional state;

909 1: transit from present functional state to Test mode.

910 The *mode* parameter enumerates specific functional behavior to be exhibited while the PHY is in Test mode.
911 It may have either of the two values.

912 0: continuous transmit;

913 1: transmit with 50% duty cycle.

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

916 0: DBPSK;

917 1: DQPSK;

918 2: D8PSK;

919 3: Not used;

920 4: DBPSK + Convolutional Code;

921 5: DQPSK + Convolutional Code;

922 6: D8PSK + Convolutional Code;

923 7: Not used.

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

926 0: Maximal output level (MOL);

927 1: MOL -3 dB;

928 2: MOL -6 dB;

929 ...

930 7: MOL -21 dB;

931 **3.11.4.8.3 Use**

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

933 **3.11.4.9 PLME_TESTMODE.confirm**

934 **3.11.4.9.1 Function**

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

937 **3.11.4.9.2 Structure**

938 The semantics of this primitive are as follows:

939 `PLME_TESTMODE.confirm{Result}.`

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

941 0: Success;

942 1: Failure. Transition to Testmode failed due to internal implementation issues;

943 2: PHY layer is already in Testmode.

944 **3.11.4.9.3 Use**

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

946 **3.11.4.10 PLME_GET.request**

947 **3.11.4.10.1 Function**

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

949 **3.11.4.10.2 Structure**

950 The semantics of this primitive is as follows:

951 `PLME_GET.request{PIBAttribute}`

952 The *PIBAttribute* parameter identifies specific attribute as enumerated in *Id* fields of tables that enumerate
953 PIB attributes (Section 6.2.2).

954 **3.11.4.10.3 Use**

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

956 **3.11.4.11 PLME_GET.confirm**

957 **3.11.4.11.1 Function**

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

960 **3.11.4.11.2 Structure**

961 The semantics of this primitive is as follows:

962 PLME_GET.confirm{status, PIBAttribute, PIBAttributeValue}

963 The *status* parameter reports the result of requested information and may have one of the values shown in
964 Table 10.

965 **Table 10 - Values of the status parameter in PLME_GET.confirm primitive**

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

966 The *PIBAttribute* parameter identifies specific attribute as enumerated in *Id* fields of tables that enumerate
967 PIB attributes (Section 6.2.2).

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

969 **3.11.4.11.3 Use**

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

971 4 MAC layer

972 4.1 Overview

973 A Subnetwork can be logically seen as a tree structure with two types of Nodes: the Base Node and Service
974 Nodes.

975 • **Base Node:** It is at the root of the tree structure and it acts as a master Node that provides all
976 Subnetwork elements with connectivity. It manages the Subnetwork resources and connections.
977 There is only one Base Node in a Subnetwork. The Base Node is initially the Subnetwork itself, and
978 any other Node should follow a Registration process to enroll itself on the Subnetwork.

979 • **Service Node:** They are either leaves or branch points of the tree structure. They are initially in a
980 Disconnected functional state and follow the Registration process in 4.6.1 to become part of the
981 Subnetwork. Service Nodes have two functions in the Subnetwork: keeping connectivity to the
982 Subnetwork for their Application layers, and switching other Nodes' data to propagate connectivity.

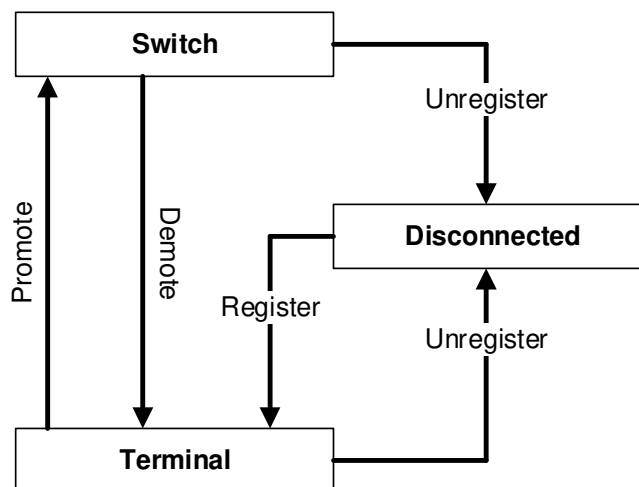
983 Devices elements that exhibit Base Node functionality continue to do so as long as they are not explicitly
984 reconfigured by mechanisms that are beyond the scope of this specification. Service Nodes, on the other
985 hand, change their behavior dynamically from "Terminal" functions to "Switch" functions and vice-versa. The
986 changing of functional states occurs in response to certain pre-defined events on the network. Figure 25
987 shows the functional state transition diagram of a Service Node.

988 The three functional states of a Service Node are **Disconnected**, **Terminal** and **Switch**:

989 • **Disconnected:** This is the initial functional state for all Service Nodes. When Disconnected, a Service
990 Node is not able to communicate data or switch other Nodes' data; its main function is to search for
991 a Subnetwork within its reach and try to register on it.

992 • **Terminal:** When in this functional state a Service Node is able to establish connections and
993 communicate data, but it is not able to switch other Nodes' data.

994 • **Switch:** When in this functional state a Service Node is able to perform all Terminal functions.
995 Additionally, it is able to forward data to and from other Nodes in the same Subnetwork. It is a branch
996 point on the tree structure.



997

998

Figure 25 - Service Node states

999 The events and associated processes that trigger changes from one functional state to another are:

- 1000 • **Registration:** the process by which a Service Node includes itself in the Base Node’s list of registered
1001 Nodes. Its successful completion means that the Service Node is part of a Subnetwork. Thus, it
1002 represents the transition between Disconnected and Terminal.
- 1003 • **Unregistration:** the process by which a Service Node removes itself from the Base Node’s list of
1004 registered Nodes. Unregistration may be initiated by either of Service Node or Base Node. A Service
1005 Node may unregister itself to find a better point of attachment i.e. change Switch Node through
1006 which it is attached to the network. A Base Node may unregister a registered Service Node as a result
1007 of failure of any of the MAC procedures. Its successful completion means that the Service Node is
1008 Disconnected and no longer part of a Subnetwork;
- 1009 • **Promotion:** the process by which a Service Node is qualified to switch (repeat, forward) data traffic
1010 from other Nodes and act as a branch point on the Subnetwork tree structure. A successful
1011 promotion represents the transition between Terminal and Switch. When a Service Node is
1012 Disconnected it cannot directly transition to Switch;
- 1013 • **Demotion:** the process by which a Service Node ceases to be a branch point on the Subnetwork tree
1014 structure. A successful demotion represents the transition between Switch and Terminal.

1015 4.2 Addressing

1016 4.2.1 General

1017 Each Node has a 48-bit universal MAC address, defined in IEEE Std 802-2001 and called EUI-48. Every EUI-48
1018 is assigned during the manufacturing process and it is used to uniquely identify a Node during the Registration
1019 process.

1020 The EUI-48 of the Base Node uniquely identifies its Subnetwork. This EUI-48 is called the Subnetwork Address
1021 (SNA).

1022 The Switch Identifier (LSID) is a unique 8-bit identifier for each Switch Node inside a Subnetwork. The
1023 Subnetwork Base Node assigns an LSID during the promotion process. A Switch Node is universally identified
1024 by the SNA and LSID. LSID = 0x00 is reserved for the Base Node. LSID = 0xFF is reserved to mean “unassigned”
1025 or “invalid” in certain specific fields (see Table 25).

1026 This special use of the 0xFF value is always made explicit when describing those fields and it shall not be used
1027 in any other field.

1028 During its Registration process, every Service Node receives a 14-bit Local Node Identifier (LNID). The LNID
1029 identifies a single Service Node among all Service Nodes that directly depend on a given Switch. The
1030 combination of a Service Node’s LNID and SID (its immediate Switch’s LSID) forms a 22-bit Node Identifier
1031 (NID). The NID identifies a single Service Node in a given Subnetwork. LNID = 0x0000 cannot be assigned to a
1032 Terminal, as it refers to its immediate Switch. LNID = 0x3FFF is reserved for broadcast and multicast traffic
1033 (see section 4.2.3 for more information). In certain specific fields, the LNID = 0x3FFF may also be used as
1034 “unassigned” or “invalid” (see Table 11 and Table 21).

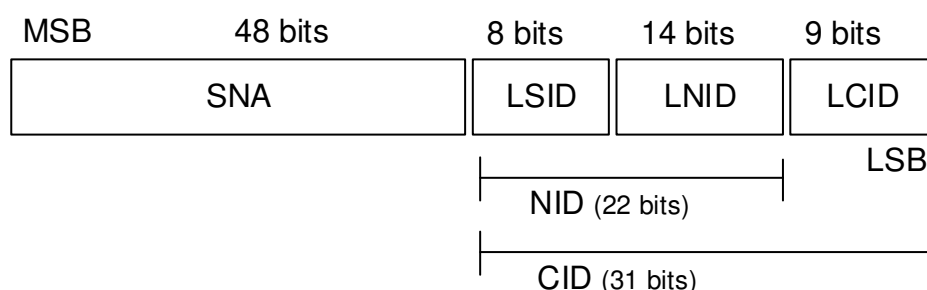
1035 This special use of the 0x3FFF value is always made explicit when describing the said fields and it shall not be
1036 used in this way in any other field.

1037 During connection establishment a 9-bit Local Connection Identifier (LCID) is reserved. The LCID identifies a
1038 single connection in a Node. The combination of NID and LCID forms a 31-bit Connection Identifier (CID). The
1039 CID identifies a single connection in a given Subnetwork. Any connection is universally identified by the SNA
1040 and CID. LCID values are allocated with the following rules:

1041 LCID=0x000 to 0x0FF, for connections requested by the Base Node. The allocation shall be made by
1042 the Base Node.

1043 LCID=0x100 to 0x1FF, for connections requested by a Service Node. The allocation shall be made by
1044 a Service Node.

1045 The full addressing structure and field lengths are shown in Figure 26



1046

1047

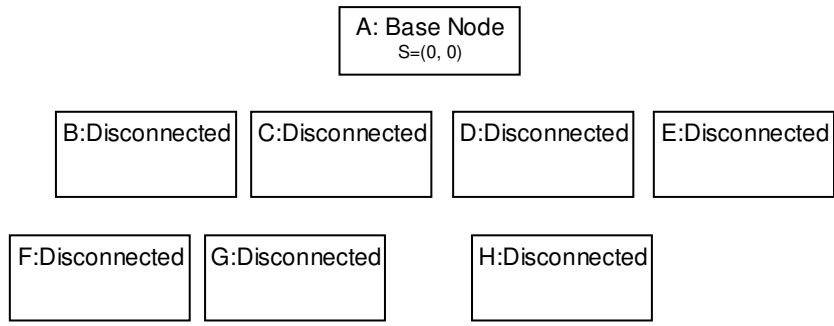
Figure 26 - Addressing Structure

1048 When a Service Node in *Terminal* state starts promotion process, the Base Node allocates a unique switch
1049 identifier which is used by this device after transition to switch state as SID of this switch. The promoted
1050 Service Node continues to use the same NID that it used before promotion i.e. it maintains SID of its next
1051 level switch for addressing all traffic generated/destined to its local application processes. To maintain
1052 distinction between the two switch identifiers, the switch identifier allocated to a Service Node during its
1053 promotion is referred to as Local Switch Identifier (LSID). Note that the LSID of a switch device will be SID of
1054 devices that connects to the Subnetwork through it.

1055 Each Service Node has a level in the topology tree structure. Service Nodes which are directly connected to
1056 the Base Node have level 0. The level of any Service Node not directly connected to the Base Node is the
1057 level of its immediate Switch plus one.

1058 4.2.2 Example of address resolution

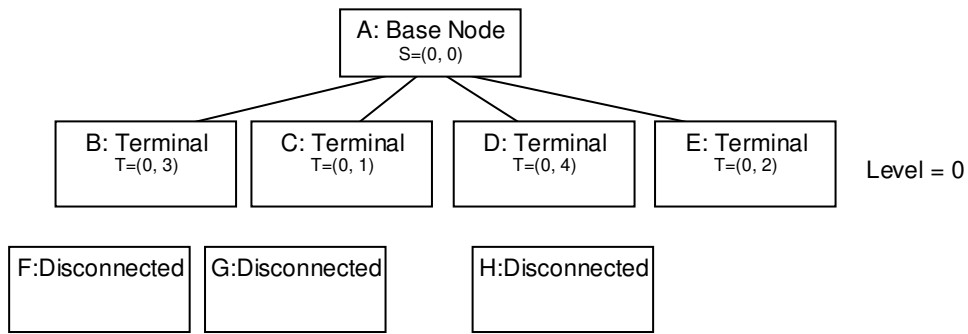
1059 Figure 27 shows an example where Disconnected Service Nodes are trying to register on the Base Node. In
1060 this example, addressing will have the following nomenclature: (SID, LNID). Initially, the only Node with an
1061 address is Base Node A, which has an NID=(0, 0).



1062
1063

Figure 27 - Example of address resolution: phase 1

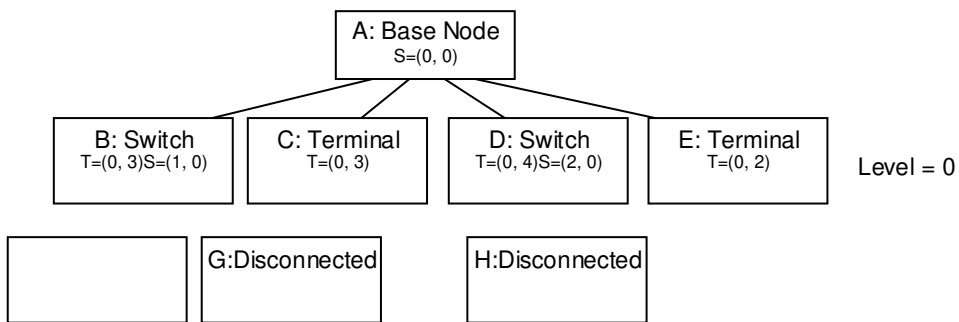
1064 Every other Node of the Subnetwork will try to register on the Base Node. Only B, C, D and E Nodes are able
1065 to register on this Subnetwork and get their NIDs. Figure 28 shows the status of Nodes after the Registration
1066 process. Since they have registered on the Base Node, they get the SID of the Base Node and a unique LNID.
1067 The level of newly registered Nodes is 0 because they are connected directly to the Base Node.



1068
1069

Figure 28 - Example of address resolution: phase 2

1070 Nodes F, G and H cannot connect directly to the Base Node, which is currently the only Switch in the
1071 Subnetwork. F, G and H will send PNPDU broadcast requests, which will result in Nodes B and D requesting
1072 promotion for themselves in order to extend the Subnetwork range. During promotion, they will both be
1073 assigned unique SIDs. Figure 29 shows the new status of the network after the promotion of Nodes B and D.
1074 Each Switch Node will still use the NID that was assigned to it during the Registration process for its own
1075 communication as a Terminal Node. The new SID shall be used for all switching functions.

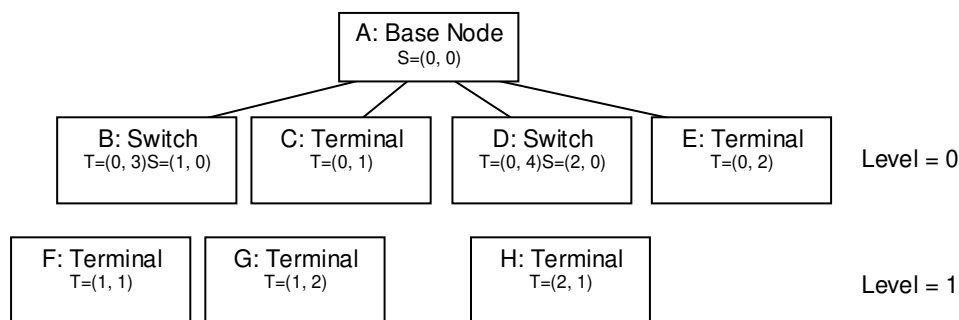


1076
1077

Figure 29 - Example of address resolution: phase 3

1078 On completion of the B and D promotion process, Nodes F, G and H shall start their Registration process and
1079 have a unique LNID assigned. Every Node on the Subnetwork will then have a unique NID to communicate
1080 like a Terminal, and Switch Nodes will have unique SIDs for switching purposes. The level of newly registered

1081 Nodes is 1 because they register with level 0 Nodes. On the completion of topology resolution and address
 1082 allocation, the example Subnetwork would be as shown in Figure 30.



1083
 1084 **Figure 30 - Example of address resolution: phase 4**

1085 4.2.3 Broadcast and multicast addressing

1086 Multicast and broadcast addresses are used for communicating data to multiple Nodes. There are several
 1087 broadcast and multicast address types, depending on the context associated with the traffic flow. Table 11
 1088 describes different broadcast and multicast addressing types and the SID and LNID fields associated with each
 1089 one.

1090 **Table 11 - 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.

1091 4.3 MAC functional description

1092 4.3.1 Service Node start-up

1093 At functional level, Service Node starts in Disconnected state. The only functions that may be performed in a
 1094 *Disconnected* functional state are: reception of any beacons on the channel and transmission of PNPDU.
 1095 Each Service Node shall maintain a Switch table that is updated with the reception of a beacon from any new
 1096 Switch Node. Based on local implementation policies, a Service Node may select any Switch Node from the
 1097 Switch table and proceed with the Registration process with that Switch Node. The criterion for selecting a
 1098 Switch Node from the Switch table is beyond the scope of this specification.

1099 Upon start, a Service Node shall operate in one of the bands in its band-plan and scan the band for at least
 1100 *macMinBandSearchTime* duration of time. On completion of this duration, the Service Node may move to

1101 next band in its band-plan or choose to continue to operate in same band for longer time. Such decisions
1102 are left to implementations.

1103 While scanning a band for available connection options, a Service Node shall listen on the band for at least
1104 *macMinSwitchSearchTime* before deciding that no beacon is being received. It may optionally add some
1105 random variation to *macMinSwitchSearchTime*, but this variation cannot be more than 10% of
1106 *macMinSwitchSearchTime*. If no beacons are received in this time, the Service Node shall broadcast a PNPDU.

1107 PNPDU's are transmitted when a Service Node is not time synchronized to an existing Subnetwork, therefore
1108 there are chances that they may collide with contention-free transmissions in a nearby Subnetwork. A Service
1109 Node shall therefore necessarily transmit PNPDU's in DBPSK_CC modulation scheme before deciding to
1110 transmit them in one of the ROBUST modulation schemes (using PHY BC frame format), if such a modulation
1111 scheme is implemented in the device. The decision making's algorithm on transitioning from one modulation
1112 scheme to other when transmitting PNPDU's and scanning for Subnetworks are left to individual
1113 implementations. So as not to flood the network with PNPDU's, especially in cases where several devices are
1114 powered up at the same time, the Disconnected Nodes shall reduce the PNPDU transmission rate when they
1115 receive PNPDU's from other sources. Disconnected Nodes shall not transmit less than one PNPDU per
1116 *macPromotionMaxTxPeriod* units of time and no more than one PNPDU per *macPromotionMinTxPeriod* units
1117 of time. The algorithm used to decide the PNPDU's transmission rate is left to the implementer.

1118 On the selection of a specific Switch Node, a Service Node shall start a Registration process by transmitting
1119 the REG control packet (4.4.2.6.3) to the Base Node. The Switch Node through which the Service Node
1120 intends to carry out its communication is indicated in the REG control packet.

1121 4.3.2 Starting and maintaining Subnetworks

1122 Base Nodes are primarily responsible for setting up and maintaining a Subnetwork. They would operate in a
1123 band comprising of one or more channels. Implementations claiming compliance with this specification shall
1124 support at least the mandatory bands required in the respective conformance specification. Base Nodes
1125 perform the following functions in order to setup and maintain a Subnetwork:

- 1126 • **Beacon transmission.** The Base Node and all Switch Nodes in the Subnetwork shall broadcast beacons
1127 at fixed intervals of time. The Base Node shall always transmit at least one beacon per super-frame.
1128 Switch Nodes shall transmit beacons with a frequency prescribed by the Base Node at the time of
1129 their promotion, which would also be at-least one beacon per super-frame.
- 1130 • **Promotion and demotion of Terminals and switches.** All promotion requests generated by Terminal
1131 Nodes upon reception of PNPDU's are directed to the Base Node. The Base Node maintains a table of
1132 all the Switch Nodes on the Subnetwork and allocates a unique SID to new incoming requests. Upon
1133 reception of multiple promotion requests, the Base Node can, at its own discretion, reject some of
1134 the requests. Likewise, the Base Node is responsible for demoting registered Switch Nodes. The
1135 demotion may either be initiated by the Base Node (based on an implementation-dependent decision
1136 process) or be requested by the Switch Node itself.
- 1137 • **Registration management.** The Base Node receives Registration requests from all new Nodes trying
1138 to be part of the Subnetwork it manages. The Base Node shall process each Registration request it
1139 receives and respond with an accept or a reject message. When the Base Node accepts the
1140 registration of a Service Node, it shall allocate an unique NID to it to be used for all subsequent

-
- 1141 communication on the Subnetwork. Likewise, the Base Node is responsible for deregistering any
1142 registered Service Node. The unregistration may be initiated by the Base Node (based on an
1143 implementation-dependent decision process) or requested by the Service Node itself.
- 1144 • **Connection setup and management:** The MAC layer specified in this document is connection-
1145 oriented, implying that data exchange is necessarily preceded by connection establishment. The Base
1146 Node is always required for all connections on the Subnetwork, either as an end point of the
1147 connection or as a facilitator (direct connections; Section 4.3.6) of the connection.
 - 1148 • **Channel access arbitration.** The usage of the channel by devices conforming to this specification may
1149 be controlled and contention-free at certain times and open and contention-based at others. The Base
1150 Node prescribes which usage mechanism shall be in force at what time and for how long.
1151 Furthermore, the Base Node shall be responsible for assigning the channel to specific devices during
1152 contention-free access periods.
 - 1153 • **Distribution of random sequence for deriving encryption keys.** When using Security Profile 1 (see
1154 4.3.8.1), all control messages in this MAC specification shall be encrypted before transmission. Besides
1155 control messages, data transfers may be optionally encrypted as well. The encryption key is derived
1156 from a 128-bit random sequence. The Base Node shall periodically generate a new random sequence
1157 and distribute it to the entire Subnetwork, thus helping to maintain the Subnetwork security
1158 infrastructure.
 - 1159 • **Multicast group management.** The Base Node shall maintain all multicast groups on the Subnetwork.
1160 This shall require the processing of all join and leave requests from any of the Service Nodes and the
1161 creation of unsolicited join and leave messages from Base Node application requests.

1162 4.3.3 Channel Access


1163 4.3.3.1 MAC Frames

1164 Time is divided into composite units of abstraction for channel usage, called MAC Frames. Composition of a
1165 MAC frame is shown in Figure 31. A frame broadly comprises of two parts:

- 1166 • **Contention Free Part (CFP):** This is the first part of a frame. Only devices that are explicitly granted
1167 permission by Base Node are allowed to transmit in CFP. Devices allocated CFP time are also given
1168 start and end time between which they need to complete their transmission and they are not allowed
1169 to use the channel for rest of the CFP duration.
- 1170 • **Shared Contention Period (SCP):** This is the second half of a frame following the CFP where devices
1171 are free to access the channel, provided they:
 - 1172 ○ Comply with CSMA CA algorithm enumerated in section 4.3.3.3.2 before transmitting their
1173 data
 - 1174 ○ Respect SCP boundaries within a MAC Frame, together with the corresponding guard-times.

1175 A guard-time of *macGuardTime* needs to be respected at both, beginning and end of CFP. Note that the
1176 length of CFP communicated in a beacon is inclusive of its respective guard-times.

1177 In order to facilitate changes to SCP and CFP times in large networks where beacons may not be transmit in
1178 every frame, a notion of super-frame is defined. A super-frame is comprised of *MACSuperFrameLength*
1179 number of frames. Each frame is numbered in modulo- *MACSuperFrameLength* manner so as to propagate
1180 information of super-frame boundary to every device in the subnetwork.



1181
1182

Figure 31 - Structure of a MAC Frame

1183 The length of a frame, *macFrameLength*, together with those of SCP and of CFP are all variable and are
1184 defined by Base Node depending on factors such as channel conditions, network size etc. The following
1185 mandatory guidelines shall be followed by Base Node implementations while defining the duration of these
1186 parameters:

- 1187 • Frame length can only be one of the four values specified for PIB attribute *macFrameLength*.
- 1188 • CFP duration within a frame shall at all times be at least $(MACBeaconLength1 + 2 \times macGuardTime)$
- 1189 • SCP duration within a frame shall at no point in time be less than *MACMinSCPLength*.

1190 Service Nodes may continue to access the channel based on frame organization communicated by Base Node
1191 in last received BPDU. Such use of channel also applies to frames when no BPDU is received by the Service
1192 Node. Non-reception of BPDU can happen either in normal course when the corresponding Switch Node does
1193 not transmit BPDU in every frame or transient channel disturbance resulting in erroneous BPDU reception.

1194 4.3.3.2 Contention-Free Period

1195 Each MAC frame shall have a contention-free period whose duration, in the least, allows transmission of one
1196 BPDU.

1197 CFP durations are allocated to Service Nodes in either of the two scenarios:

- 1198 • As part of promotion procedure carried out for a Terminal node. In all such cases, the CFP allocation
1199 will be for usage as beacon-slot by the Service Node being promoted.
- 1200
- 1201 • As part of CFP allocation process that could be initiated either from Base Node or Service Node, for
1202 use to transport application data.

1203 Service Nodes make channel allocation request in a CFP MAC control packet. The Base Node acts on this
1204 request and responds with a request acceptance or denial. In the case of request acceptance, the Base Node
1205 shall respond with the location of allocation time within MAC frame, the length of allocation time and number
1206 of future MAC frames from which the allocation pattern will take effect. The allocation pattern remains
1207 effective unless there is an unsolicited location change of the allocation period from the Base Node (as
1208 a result of channel allocation pattern reorganization) or the requesting Service Node sends an explicit de-
1209 allocation request using a CFP MAC control packet.

1210 Changes resulting from action taken on a CFP MAC control message that impact overall MAC frame structure
1211 are broadcast to all devices using an FRA MAC control message.

1212 All CFP_ALC_REQ_S requests coming from Terminal or Switch Nodes are addressed to the Base Node.
1213 Intermediate Switch Nodes along the transmission path merely act on the allocation decision by the Base
1214 Node.

1215 Base Nodes may allocate overlapping times to multiple requesting Service Nodes. Such allocations may lead
1216 to potential interference. Thus, a Base Node should make such allocations only when devices that are
1217 allocated channel access for concurrent usage are sufficiently separated. In a multi-level Subnetwork, when
1218 a Service Node that is not directly connected to the Base Node makes a request for CFP, the Base Node shall
1219 allocate CFPs to all intermediate Switch Nodes such that the entire transit path from the source Service Node
1220 to Base Node has contention-free time-slots reserved. The Base Node shall transmit multiple CFP control
1221 packets. The first of these CFP_ALC_IND will be for the requesting Service Node. Each of the rest will be
1222 addressed to an intermediate Switch Node.

1223 **4.3.3.2.1 Beacons**

1224 **4.3.3.2.1.1 General**

1225 Base Node and every other Switch Node in a Subnetwork transmit a Beacon PDU (BPDU) at least once per
1226 super-frame. A BPDU contains administrative and operational information of its respective Subnetwork. Its
1227 contents are enumerated in 4.4.4. Every Service Node in a Subnetwork is required to track beacons as
1228 explained in 4.3.4.1. In addition to using the administrative and operational information, Service Nodes also
1229 synchronize their notion of time based on time of reception of BPDUs.

1230 Since BPDUs are important to keep a Subnetwork running, Base Node and every Switch Node transmitting a
1231 BPDU shall do so using a robust modulation scheme, which is either DBPSK_CC, DBPSK_R or DQPSK_R.
1232 Beacons in DBPSK_CC are transmitted using PHY Frame Type A. Beacons in DBPSK_R and DQPSK_R are
1233 transmitted in PHY Frame Type B. Note that the chosen modulation scheme shall be compliant with the
1234 definition of *macRobustnessManagement*. Every device, including the Base Node shall transmit BPDU with
1235 maximum output power.

1236 **4.3.3.2.1.2 Beacon-slots**

1237 Unit of time dedicated for transmission of a BPDU is called a beacon-slot. Depending on the corresponding
1238 modulation it has a length of either $(MACBeaconLength1 + macGuardTime)$, $(MACBeaconLength2 +$
1239 $macGuardTime)$ or $(MACBeaconLength3 + macGuardTime)$. Note that it includes not only the time required
1240 to transmit the BPDU but also the *macGuardTime* that is required to ensure minimal separation between
1241 successive transmissions.

1242 Note that a Base Node:

- 1243 • May decide to not use its beacon-slot in every frame implying that it transmits BPDU at less than once
1244 per frame frequency
- 1245 • Will necessarily use its beacon-slot at least once per *MACSuperFrameLength*

-
- 1246 • May allocate its beacon-slot location to other switches in its network for frames where it decides to
1247 not transmit its BPDU.

1248 Every Switch Node in the Subnetwork needs to have a beacon-slot allocated in order for it to transmit its
1249 BPDU. Switch Nodes are allocated a beacon-slot at time of their promotion by the Base Node.

- 1250 • Beacon-slot allocations shall necessarily be contained within the CFP duration of a frame.
1251 • Base Node may time-multiplex beacon-slots i.e. allocate same duration of time to different switches
1252 in different frames.
1253 • A Switch Node may request to change the duration of its beacon-slot when it decides to change the
1254 modulation scheme of its BPDU.

1255 With the Registration of each new Switch on the Subnetwork, the Base Node may change the modulation,
1256 beacon-slot or BPDU transmission frequency (or both) of already registered Switch devices. When such a
1257 change occurs, the Base Node transmits an unsolicited PRO_REQ to each individual Switch device that is
1258 affected. The Switch device addressed in the PRO_REQ shall transmit an acknowledgement, PRO_ACK, back
1259 to the Base Node. During the reorganization of beacon-slots, if there is a change in CFP duration, the Base
1260 Node shall transmit an FRA control packet to the entire Subnetwork. The BN also sends a FRA control packet
1261 in advance of a change in length of a frame.

1262 Switch devices that receive an FRA control packet shall relay it to their entire control domain because FRA
1263 packets are broadcast information about changes to frame structures.

1264 This is required for the entire Subnetwork to have a common understanding of frame structure, especially in
1265 regions where the controlling Switch devices transmit BPDUs at frequencies below once per frame.

1266 **4.3.3.3 Shared-contention period**

1267 **4.3.3.3.1 General**

1268 Shared-contention period (SCP) is the time when any device in Subnetwork can transmit data. SCP follows
1269 the CFP duration within a frame and its duration is defined by Base Node. Collisions resulting from
1270 simultaneous attempt to access the channel are avoided by the CSMA-CA mechanism specified in this section.
1271 SCP durations are highlighted by the following key specifications:

- 1272 • SCP duration within a frame shall at no point in time be less than *MACMinSCPLength*.
1273 • Maximum possible duration of SCP shall be $(macFrameLength - (MACBeaconLength1 + 2 \times$
1274 *macGuardTime*)). This is the case of a subnetwork that does not have dedicated CFP requests from
1275 any Service Node.

1276 **4.3.3.3.2 CSMA-CA algorithm**

1277 The CSMA-CA algorithm implemented in devices works as shown in Figure 32.

1278 Implementations start with a random backoff time (*macSCPRBO*) based on the priority of data queued for
1279 transmission. *MACPriorityLevels* levels of priority need to be defined in each implementation, with a lower
1280 value indicating higher priority. In the case of data aggregation, the priority of aggregate bulk is governed by
1281 the highest priority data it contains. The *macSCPRBO* for a transmission attempt is given as below:

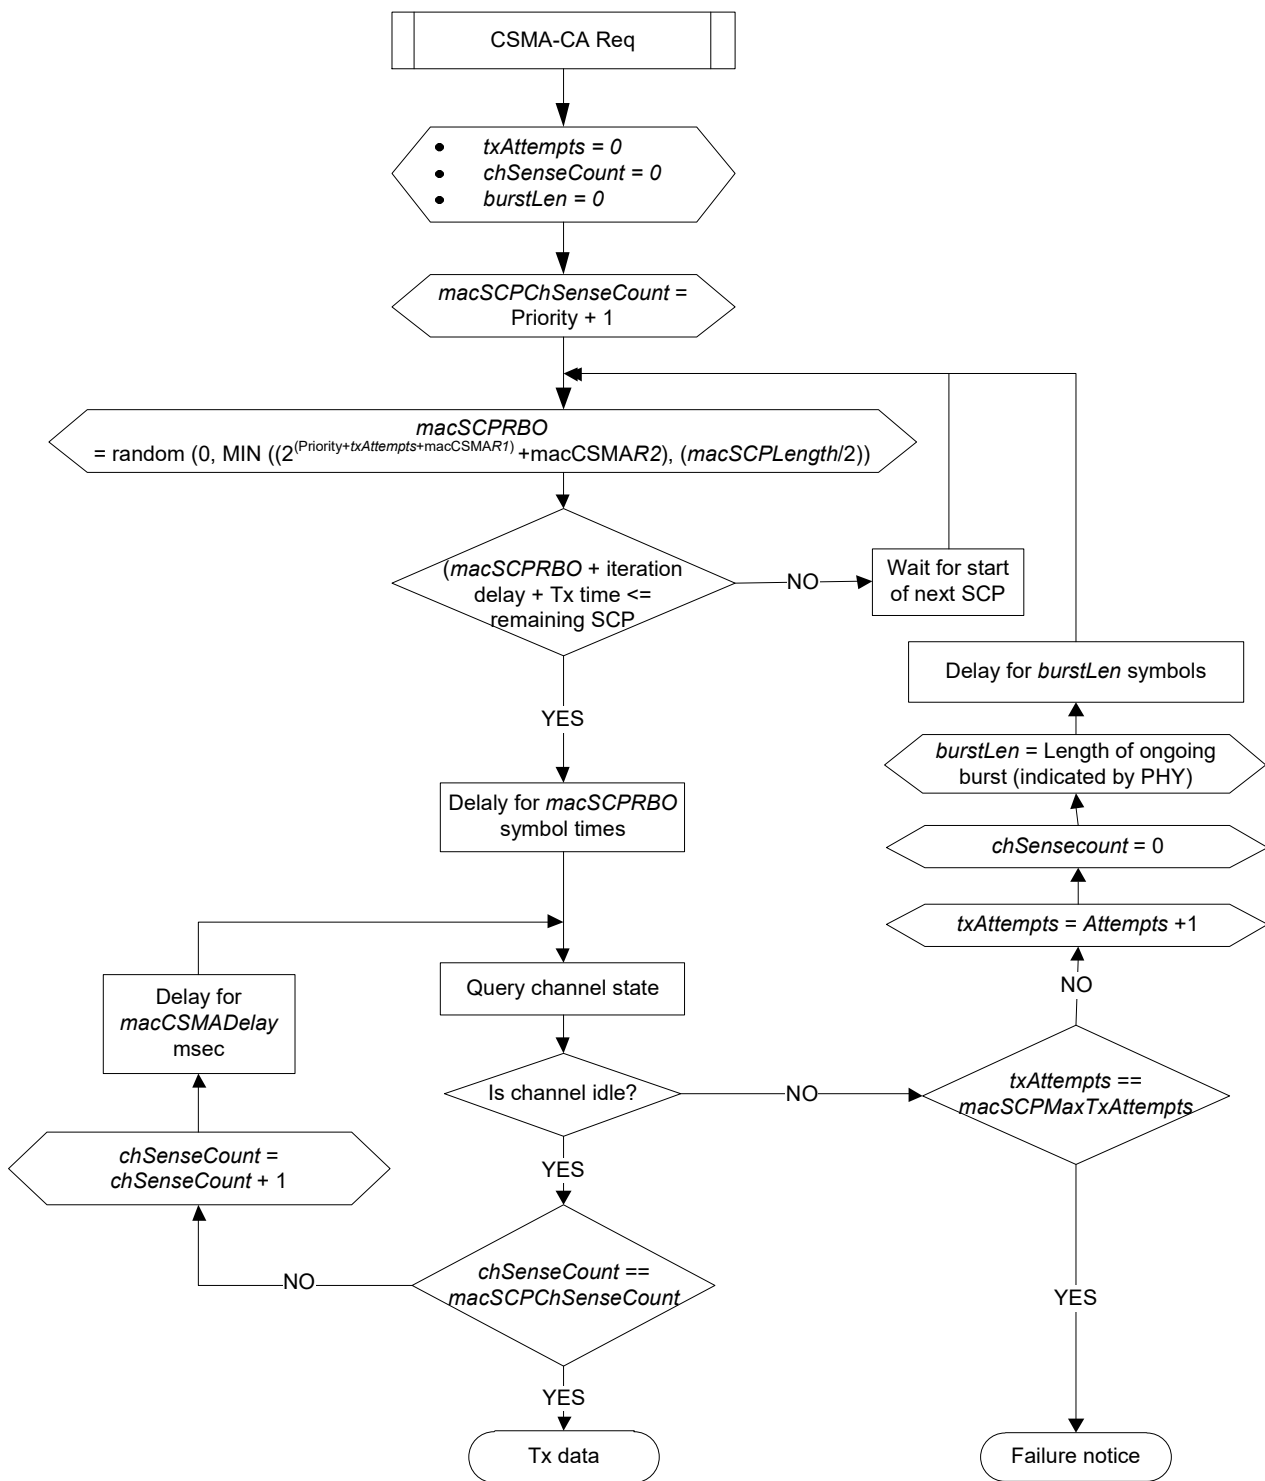
1282 $macSCPRBO = \text{random}(0, \text{MIN}((2^{(\text{Priority} + \text{txAttempts} + \text{macCSMAR1})} + \text{macCSMAR2}), (\text{macSCPLength}/2)))$

1283 or when Robust Modes are supported:

1284 $macSCPRBO = \text{random}(0, \text{MIN}((2^{(\text{Priority} + \text{txAttempts} + \text{macCSMAR1Robust})} + \text{macCSMAR2Robust}), (\text{macSCPLength}/2)))$

1285 *macCSMAR1/macCSMAR1Robust* and *macCSMAR2/macCSMAR2Robust* control the initial contention
1286 window size. *macCSMAR1/macCSMAR1Robust* helps to increase the contention window size exponentially
1287 while *macCSMAR2/macCSMAR2Robust* helps to increase the contention window linearly. A higher value of
1288 *macCSMAR1/macCSMAR1Robust* and/or *macCSMAR2/macCSMAR2Robust* is recommended for large
1289 networks. It is recommended to not decrease the default values.

1290 Before a backoff period starts, a device should ensure that the remaining SCP time is long enough to
1291 accommodate the backoff, the number of iterations for channel-sensing (based on data priority) and the
1292 subsequent data transmission. If this is not the case, backoff should be aborted till the SCP starts in the next
1293 frame. Aborted backoffs that start in a subsequent frame should not carry *macSCPRBO* values of earlier
1294 attempts. *macSCPRBO* values should be regenerated on the resumption of the transmission attempt in the
1295 SCP time of the next frame.



1296

1297

Figure 32 - Flow chart for CSMA-CA algorithm

1298

1299

1300

1301

On the completion of *macSCRBO* symbol time, implementations perform channel-sensing. Channel sensing shall be performed one or more times depending on 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:

1302

$$macSCPChSenseCount = Priority + 1$$

1303 and each channel sense should be separated by a *macCSMADelay* ms delay.

1304 **Note:** *macSCPRBO* and *macCSMADelay* follow a different range and different default value depending on the
1305 modulation scheme that is intended to be used for a transmission burst. If a device intends to use robust mode
1306 for some bursts, the values are conservative to account for extended PHY Frame (Type B) timings. The
1307 applicable values are listed in **6.2.3.2**. Implementations shall conform to listed range and default value
1308 corresponding to the modulation scheme used.

1309 When a channel is sensed to be idle on all *macSCPChSenseCount* occasions, an implementation may conclude
1310 that the channel status is idle and carry out its transmission immediately.

1311 During any of the *macSCPChSenseCount* channel-sensing iterations, if the channel is sensed to be occupied,
1312 implementations should reset all working variables. The local counter tracking the number of times a channel
1313 is found to be busy should be incremented by one and the CSMA-CA process should restart by generating a
1314 new *macSCPRBO*. The remaining steps, starting with the backoff, should follow as above.

1315 If the CSMA-CA algorithm restarts *macSCPMaxTxAttempts* number of times due to ongoing transmissions
1316 from other devices on the channel, the transmission shall abort by informing the upper layers of CSMA-CA
1317 failure.

1318 **4.3.3.3 MAC control packet transmission**

1319 MAC control packets (0) shall follow the following channel access rules:

- 1320 • Always transmit in SCP.
- 1321 • Use priority level of *MACCtrlPktPriority*.
- 1322 • The MAC Control Packets shall be transmitted in a modulation scheme robust enough to reach the
1323 receiving peer but no less robust than DBPSK_CC.
- 1324 • Transmitted with PHY Frame Type B for DBPSK_R and DQPSK_R. For all other modulation schemes,
1325 control packets are transmitted using PHY Frame Type A.

1326 **4.3.4 Tracking switches and peers**

1327 **4.3.4.1 Tracking switches**

1328 Service Nodes shall keep track of all neighboring Switch Nodes by maintaining a list of beacons received. Such
1329 tracking shall keep a Service Node updated on reception signal quality from Switch Nodes other than the one
1330 to which it is connected, thus making it possible to change connection points (Switch Node) to the
1331 Subnetwork if link quality to the existing point of connectivity degrades beyond an acceptable level.

1332 Note that such a change of point of connectivity may be complex for Switch Nodes because of devices
1333 connected through them. However, at certain times, network dynamics may justify a complex reorganization
1334 rather than continue with existing limiting conditions.

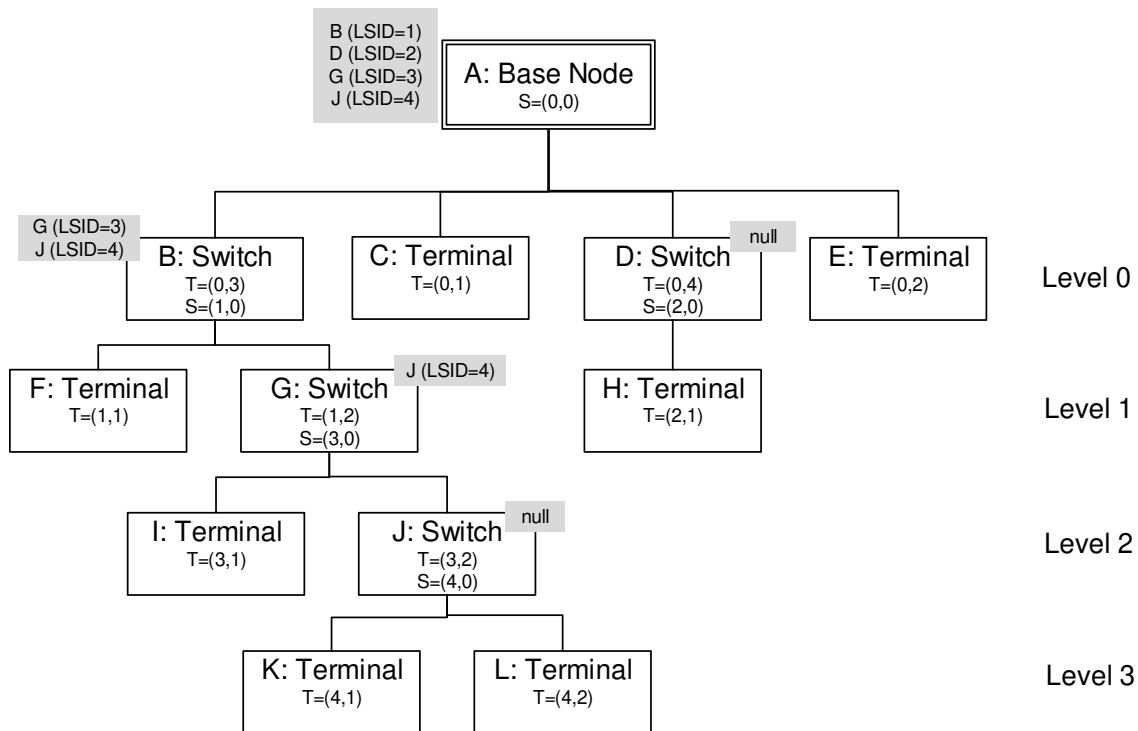
1335 **4.3.4.2 Tracking disconnected Nodes**

1336 Terminals shall process all received PNPDU. When a Service Node is Disconnected, it doesn't have
 1337 information on current MAC frame structure so the PNPDU may not necessarily arrive during the SCP. Thus,
 1338 Terminals shall also keep track of PNPDU during the CFP or beacon-slots.

1339 On processing a received PNPDU, a Terminal Node may decide to ignore it and not generate any
 1340 corresponding promotion request (PRO_REQ_S). Receiving multiple PNPDU can indicate that there is no
 1341 other device in the vicinity of Disconnected Nodes, implying that there will be no possibility of new devices
 1342 for connecting to the Subnetwork if the Terminal Node does not request promotion itself. A Terminal Node
 1343 shall ignore no more than *MACMaxPRNIgnore* PNPDU. After this maximum number of ignored PNPDU the
 1344 Terminal Node shall start a Promotion procedure as described in 4.6.3. The time in which the procedure will
 1345 start shall be randomly selected in the range of $[0, MACMaxPRNIgnore * macPromotionMinTxPeriod]$ seconds.

1346 **4.3.4.3 Tracking switches under one node**

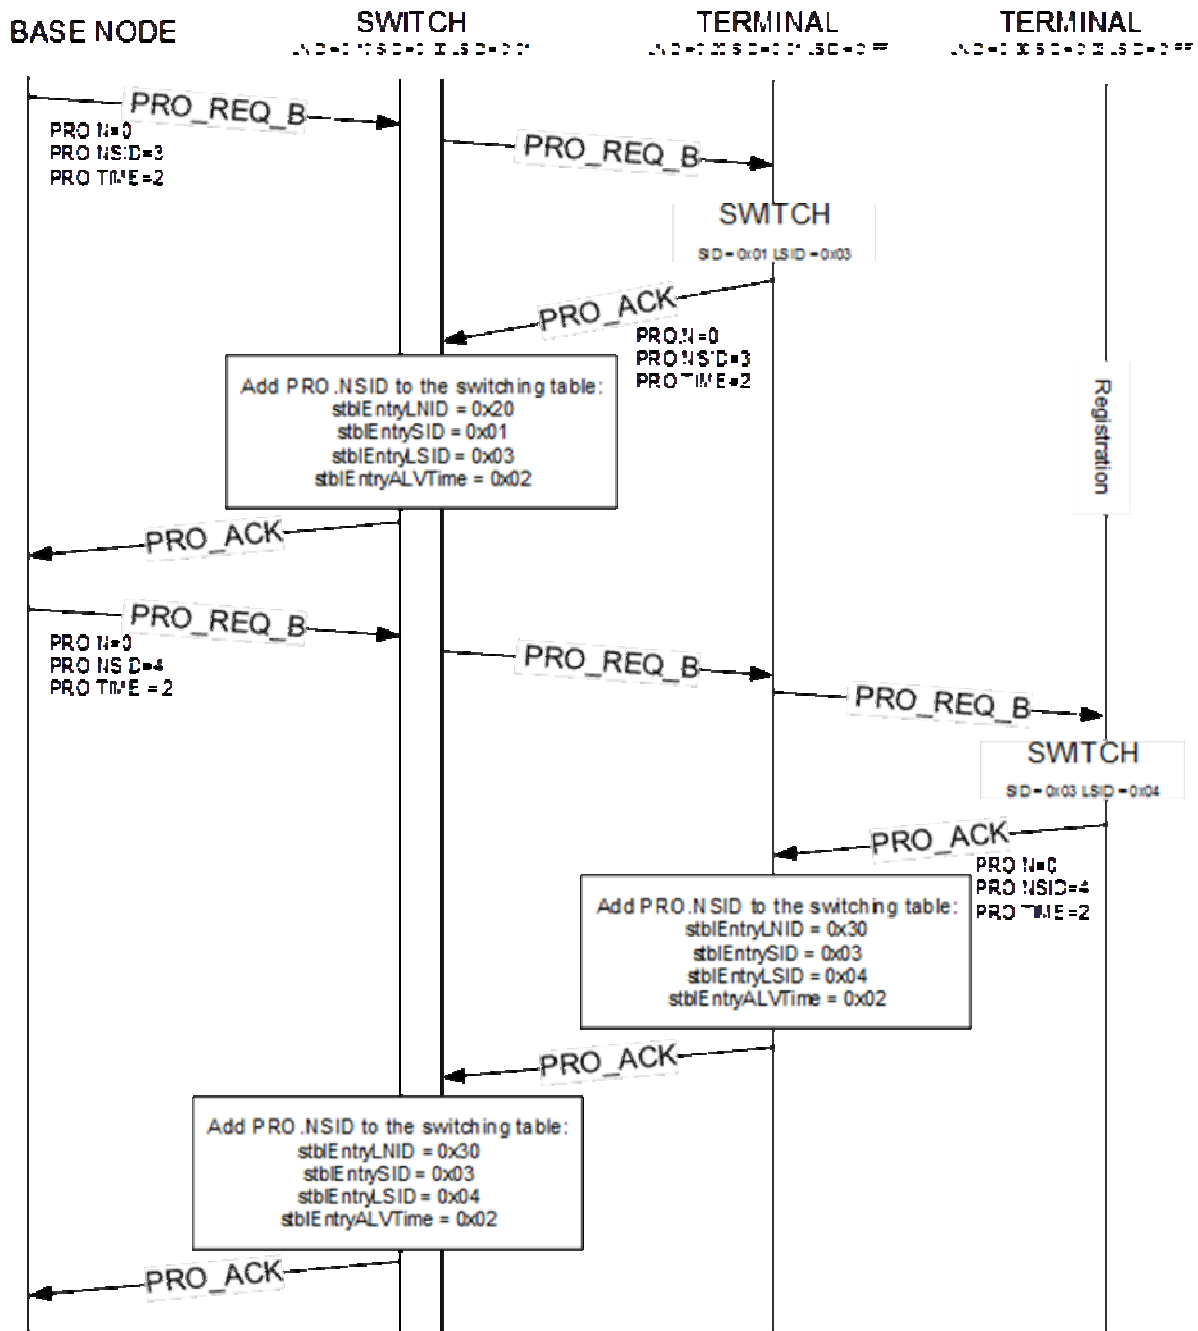
1347 Service Nodes in *Switch* functional state shall keep track of the Switches under their tree by maintaining the
 1348 *maListSwitchTable*. Maintaining this information is sufficient for switching because traffic to/from Terminal
 1349 Nodes will also contain the identity of their respective Switch Nodes (PKT.SID). Thus, the switching function
 1350 is simplified in that maintaining an exhaustive listing of all Terminal Nodes connected through it is not
 1351 necessary. After promotion Switch Nodes start with no entries in their switching table.



1352 **Figure 33 -Switching tables examples**

1353 One Switch Node shall include in the switching list the SID of every promoted Terminal node which is directly
 1354 connected to it or to one Switch already included in the *maListSwitchTable*. In this case the Node shall create
 1355 an entry with the *stblEntryLSID* value equals to the NSID field of the PRO_ACK packet, the *stblEntrySID* value
 1356

1357 equal to PKT.SID of the PRO.ACK Packet Header and stblEntryLNID equal to PKT.LNID of the PRO.ACK Packet
 1358 Header.



1359
 1360 **Figure 34 - Filling example for the switching table for Switch of Level 0.**

1361 Similarly the Switch Node shall mark the entry to be removed from the list the SID when a node is removed
 1362 or unregistered. This is be done by listening the PRO_DEM_B, PRO_DEM_S, REG_UNR_B or REG_UNR_S
 1363 packets.

1364 Each entry of the *macListSwitchTable* also contains the information related to the Alive time related to the
 1365 Switch node. The stblEntryALVTime is updated with the TIME field received during the promotion, beacon
 1366 robustness change or the keep alive procedures. The Switch Node shall also maintain the T_{keep-alive} timer for

1367 every Switch under its tree. The Switch Node shall refresh the timers as specified in section 4.6.5. If the T_{keep-}
1368 $alive$ timer expires the entry in the *macListSwitchTable* shall be marked to be removed.

1369 Every time an entry is marked to be removed, the Switch Node shall check if the *stblEntrySID* of other entries
1370 is equal to the *stblEntryLSID*. In these cases all the entries shall be marked to be removed, meaning that one
1371 entire branch has left the network.

1372 When one entry is marked to be removed the Switch Node shall wait (*macCtrlMsgFailTime* +
1373 *macMinCtrlReTxTimer*) seconds. This time ensures that all retransmit packets which use the SID have left the
1374 Subnetwork. When the timer expires the table entry shall be removed.

1375 **4.3.5 Switching**

1376 **4.3.5.1 General**

1377 On a Subnetwork, the Base Node cannot communicate with every Node directly. Switch Nodes relay traffic
1378 to/from the Base Node so that every Node on the Subnetwork is effectively able to communicate with the
1379 Base Node. Switch Nodes selectively forward traffic that originates from or is destined to one of the Service
1380 Nodes in its control hierarchy. All other traffic is discarded by Switches, thus optimizing traffic flow on the
1381 network.

1382 Different names of MAC header and packets are used in this section. Please refer to the section 4.4.2 to find
1383 their complete specification.

1384 **4.3.5.2 Switching process**

1385 Switch Nodes forward traffic to their control domain in a selective manner. The received data shall fulfill the
1386 conditions listed below for it to be switched. If the conditions are not met, the data shall be silently discarded.

1387 Downlink packets (*HDR.DO*=1) shall meet any of the following conditions in order to be switched:

- 1388 • Destination Node of the packet is connected to the Subnetwork through this Switch Node, i.e. *PKT.SID*
1389 is equal to this Switch Node's *SID* or its switching table contains an entry for *PKT.SID*.
- 1390 • The packet has broadcast destination (*PKT.LNID* = 0x3FFF) and was sent by the Switch this Node is
1391 registered through (*PKT.SID*=*SID* of this Switch Node).
- 1392 • The packet has a multicast destination (*PKT.LNID*=0x3FFE), it was sent by the Switch this Node is
1393 registered through (*PKT.SID*=*SID* of this Switch Node) and at least one of the Service Nodes connected
1394 to the Subnetwork through this Switch Node is a member of the said multicast group, i.e. *LCID*
1395 specifies a group that is requested by any downstream Node in its hierarchy.

1396 Uplink packets (*HDR.DO*=0) shall meet either of the following conditions in order to be switched:

- 1397 • The packet source Node is connected to the Subnetwork through this Switch Node, i.e. *PKT.SID* is
1398 equal to this Switch Node's *SID* or its switching table contains an entry for *PKT.SID*.
- 1399 • The packet has a broadcast or multicast destination (*PKT.LNID* = 0x3FFF or 0x3FFE) and was
1400 transmitted by a Node connected to the Subnetwork through this Switch Node i.e. *PKT.SID* is equal to
1401 this Switch Node's *SID* or its switching table contains an entry for *PKT.SID*.

1402 If a packet meets previous conditions, it shall be switched. For unicast packets, the only operation to be
1403 performed during switching is to queue it to be resent in a MAC PDU with the same HDR.DO.

1404 In case of broadcast or multicast packets, the PKT.SID must be replaced with the Switch Node's LSID for
1405 Downlink packets. Uplink packets shall be resent unchanged.

1406 **4.3.5.3 Switching of broadcast packets**

1407 The switching of broadcast MAC frames operates in a different manner to the switching of unicast MAC
1408 frames. Broadcast MAC frames are identified by PKT.LNID=0x3FFF.

1409 When HDR.DO=0, i.e. the packet is an uplink packet, it is unicast to the Base Node. A Switch which receives
1410 such a packet shall apply the scope rules to ensure that it comes from a lower level and, if so, Switch it
1411 upwards towards the base. The rules given in section 4.3.5.2 must be applied. The same modulation scheme
1412 and output power level as used for unicast uplink switching shall be used.

1413 When HDR.DO=1, i.e. the packet is a Downlink packet, it is broadcast to the next level. A Switch which
1414 receives such a packet shall apply the scope rules to ensure that it comes from the higher level and, if so,
1415 switch it further to its Subnetwork. The same modulation scheme as used for beacon transmission at the
1416 maximum output power level implemented in the device shall be used so that all the devices directly
1417 connected to the Switch Node can receive the packet. The rules given in section 4.3.5.2 must be applied. The
1418 Service Node shall also pass the packet up to its MAC SAP to applications which have registered to receive
1419 broadcast packets using the MAC_JOIN service.

1420 When the Base Node receives a broadcast packet with HDR.DO=0, it shall pass the packet up its MAC SAP to
1421 applications which have registered to receive broadcast packets. The Base Node shall also transmit the packet
1422 as a Downlink packet, i.e. HDR.DO=1, using the same modulation scheme as used for beacon transmission at
1423 the maximum output power level and following the rules given in section 4.3.5.2.

1424 **4.3.5.4 Switching of multicast packets**

1425 Switch Nodes shall maintain a multicast switching table. This table contains a list of multicast group LCIDs
1426 that have members connected to the Subnetwork through the Switch Node. The LCID of multicast traffic in
1427 both Downlink and uplink directions is checked for a matching entry in the multicast switching table.
1428 Multicast traffic is only switched if an entry corresponding to the LCID is available in the table; otherwise, the
1429 traffic is silently discarded.

1430 A multicast switching table is established and managed by examining the multicast join messages (MUL
1431 control packet) which pass through the Switch. On a successful group join from a Service Node in its control
1432 hierarchy, a Switch Node adds a new multicast Switch entry for the group LCID, where necessary. An entry
1433 from the multicast switching table can be removed by the Base Node using the multicast leave procedure
1434 (see section 4.6.7.4.2). All entries from the multicast switching table shall be removed when a switch is
1435 demoted or unregistered. The multicast packet switching process depends on the packet direction.

1436 When HDR.DO=0 and PKT.LNID=0x3FFE, i.e. the packet is an uplink multicast packet, it is unicast towards the
1437 Base Node. A Switch Node that receives such a packet shall apply the scope rules to ensure it comes from a
1438 lower hierarchical level and, if so, switch it upwards towards the Base Node. No LCID-based filtering is

1439 performed. All multicast packets are switched, regardless of any multicast Switch entries for the LCID. The
1440 rules given in section 4.3.5.2 must be applied. The same modulation scheme and output power level as used
1441 for unicast uplink switching shall be used.

1442 When HDR.DO=1 and PKT.LNID=0x3FFE, i.e. the packet is a Downlink multicast packet, the multicast
1443 switching table is used. If there is an entry with the LCID corresponding to PKT.LCID in the packet, the packet
1444 is switched downwards to the part of Subnetwork controlled by this switch. The multicast traffic shall be
1445 relayed using a modulation scheme which is robust enough to ensure that all direct children which are part
1446 of the multicast group or which need to switch the multicast traffic can receive the packet. As a guideline,
1447 the same modulation scheme as used for beacon transmission at the maximum output power level can be
1448 used. The rules given in section 4.3.5.2 shall be applied. If the Service Node is also a member of the multicast
1449 group, it shall also pass the packet up its MAC SAP to applications which have registered to receive the
1450 multicast packets for that group.

1451 When the Base Node receives a multicast packet with HDR.DO=0 and it is a member of the multicast group,
1452 it shall pass the packet up its MAC SAP to applications which have registered to receive multicast packets for
1453 that group. The Base Node shall switch the multicast packet if there is an appropriate entry in its multicast
1454 switching table for the LCID, transmitting the packet as a Downlink packet, i.e. HDR.DO=1. To transmit a
1455 downlink multicast packet by the Base Node the same rules apply as for transmitting a downlink multicast
1456 packet by a switch.

1457 **4.3.6 Direct connections**

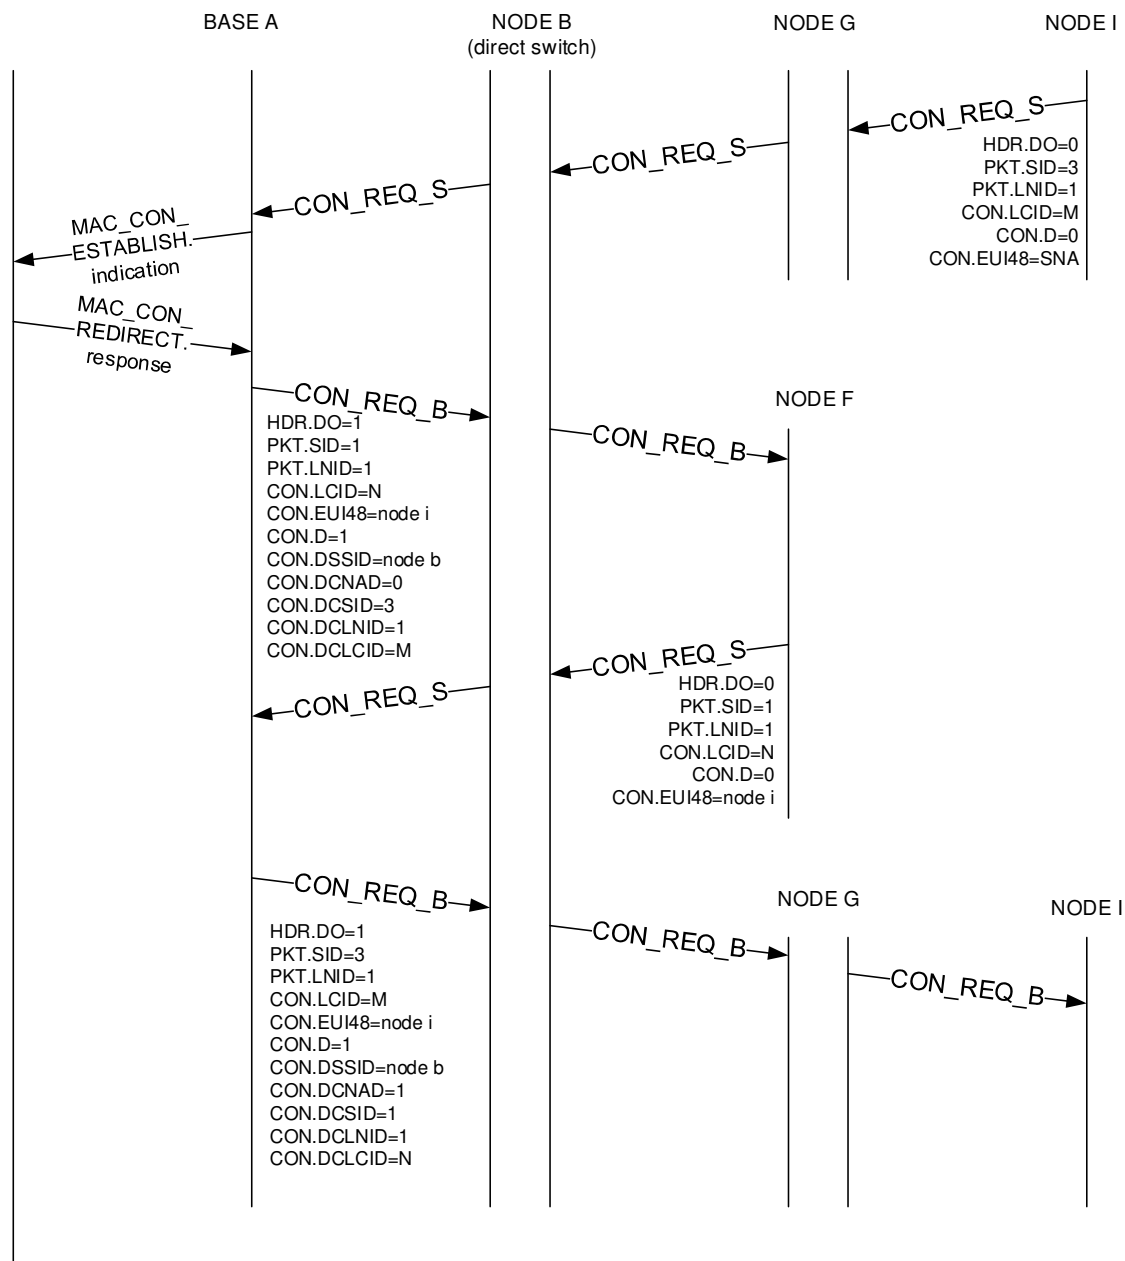
1458 **4.3.6.1 Direct connection establishment**

1459 The direct connection establishment is a little different from a normal connection although the same packets
1460 and processes are used. It is different because the initial connection request may not be acknowledged until
1461 it is already acknowledged by the target Node. It is also different because the CON_REQ_B packets shall carry
1462 information for the “direct Switch” to update the “direct switching table”.

1463 A direct switch is not different than a general switch. It is only a logical distinction of identifying the first
1464 common switch between two service-nodes that need to communicate with each other. Note that in absence
1465 of such a common switch, the Base Node would be the direct switch.

1466 There are two different scenarios for using directed connections. These scenarios use the network shown in
1467 Figure 35.

1468 The first is when the source Node does not know the destination Service Node’s EUI-48 address. The Service
1469 Node initiates a connection to the Base Node and the Base Node Convergence layer redirects the connection
1470 to the correct Service Node.



1471

1472

Figure 35 - Directed Connection to an unknown Service Node

1473 The steps to establish a direct connection, as shown in Figure 35, shall be:

- 1474 • When Node I tries to establish connection with Node F, it shall send a normal connection request
1475 (CON_REQ_S).
- 1476 • Then, due to the fact that the Base Node knows that F is the target Service Node, it should send a
1477 connection request to F (CON_REQ_B). This packet will carry information for direct Switch B to include
1478 the connection in its direct switching table.
- 1479 • F may accept the connection. (CON_REQ_S).
- 1480 • Now that the connection with F is fully established, the Base Node will accept the connection with I
1481 (CON_REQ_B). This packet will carry information for the direct Switch B to include in its direct
1482 switching table.

1483 After finishing this connection-establishment process, the direct Switch (Node B) should contain a direct
1484 switching table with the entries shown in Table 12.

1485 **Table 12 - 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

1486 The direct switching table should be updated every time a Switch receives a control packet that meets the
1487 following requirements.

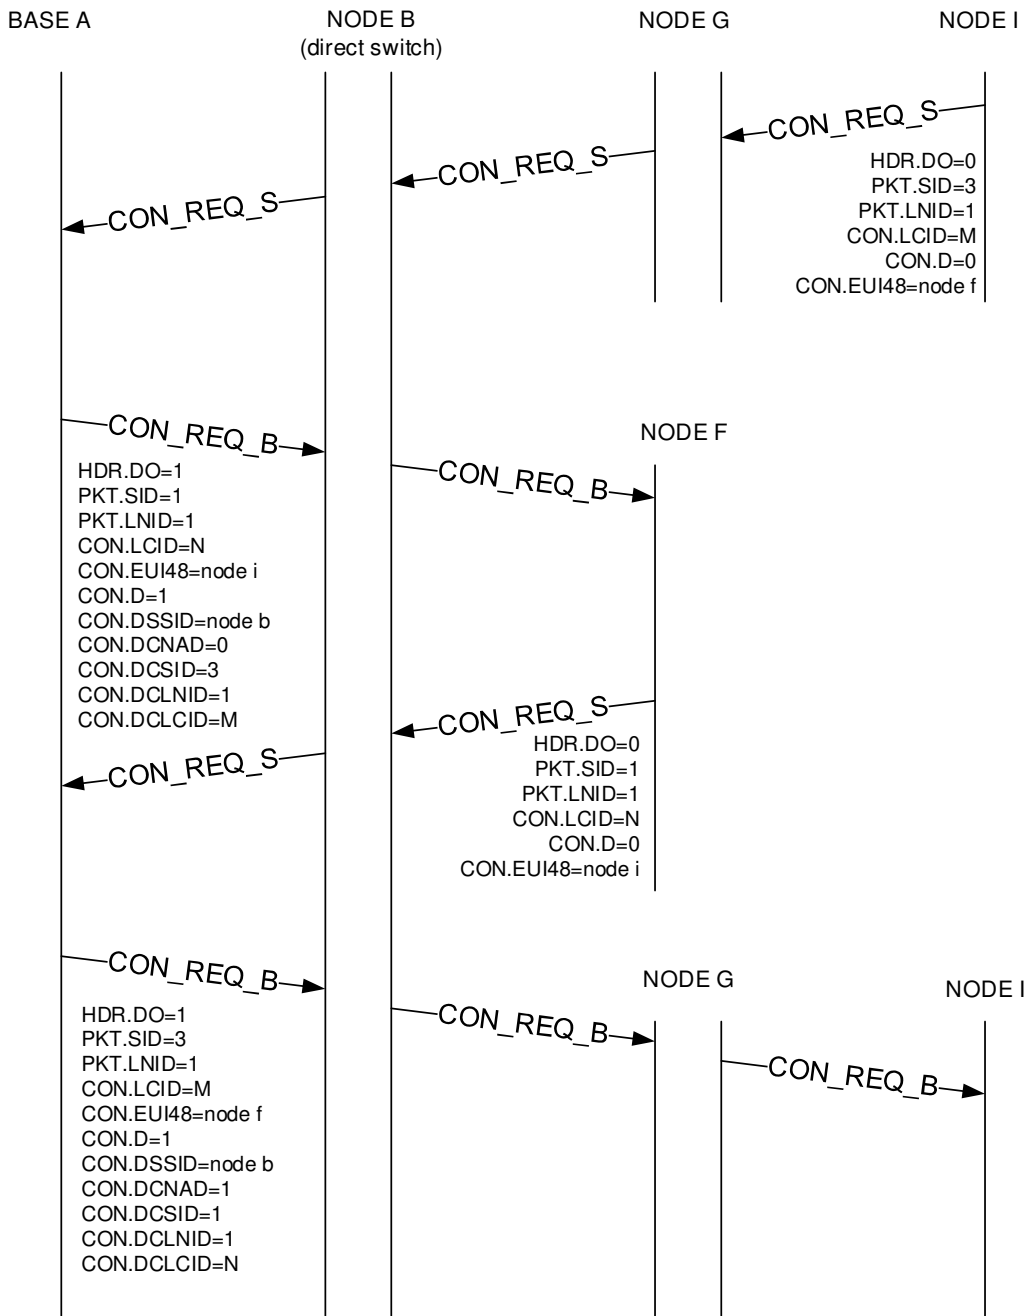
- 1488 • It is CON_REQ_B packet: HDR.DO=1, CON.TYPE=1 and CON.N=0;
- 1489 • It contains “direct” information: CON.D=1;
- 1490 • The direct information is for itself: CON.DSSID is the SID of the Switch itself.

1491 Then, the direct switching table is updated with the information:

- 1492 • Uplink (SID, LNID, LCID) = (PKT.SID, PKT.LNID, CON.LCID);
- 1493 • Downlink (SID, LNID, LCID, NAD) = (CON.DCSID, CON.DCLNID, CON.DCLCID, CON.DCNAD).

1494 The connection closing packets should be used to remove the entries.

1495 The second scenario for using directed connections is when the initiating Service Node already knows the
1496 destination Service Node’s EUI-48 address. In this case, rather than using the Base Node’s address, it uses
1497 the Service Node’s address. In this case, the Base Node Convergence layer is not involved. The Base Node
1498 MAC layer connects Service Node I directly to Service Node F. The resulting Switch table entries are identical
1499 to the previous example. The exchange of signals is shown in Figure 36.



1501

1502

Figure 36 - Example of direct connection: connection establishment to a known Service Node

1503

4.3.6.2 Direct connection release

1504

The release of a direct connection is shown in Figure 37. The signaling 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.

1505

1506

1507

1508

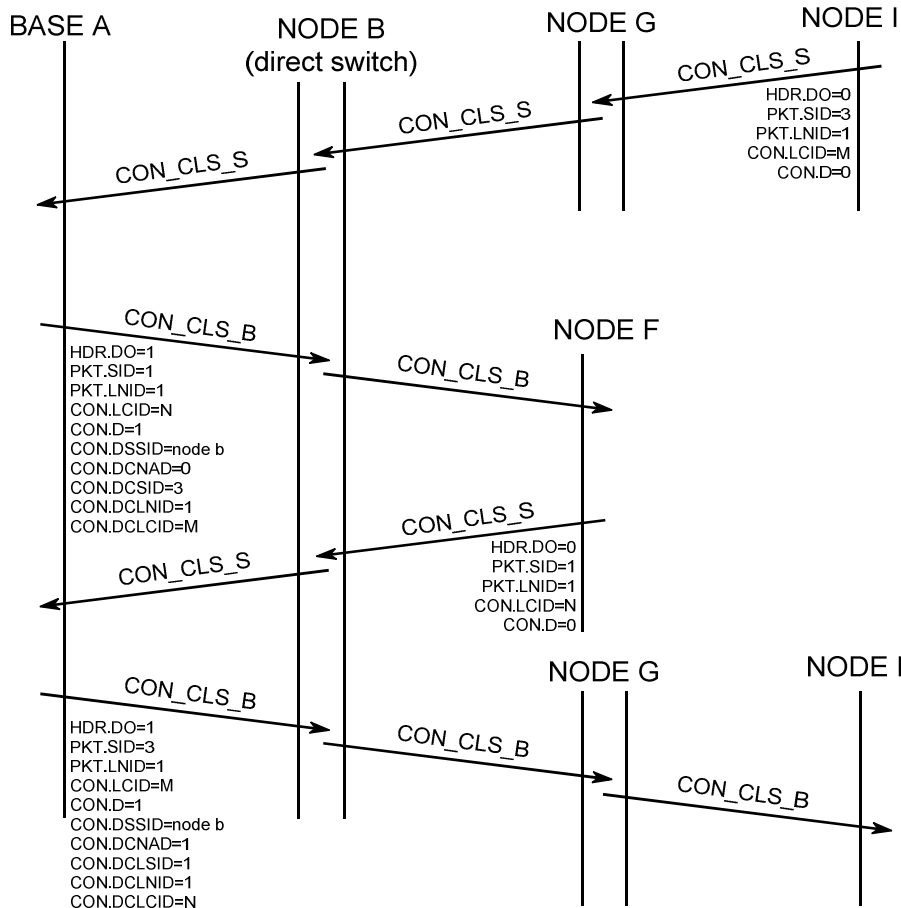
- It is `CON_CLOSE_B` packet: HDR.DO=1, CON.TYPE=1 and CON.N=1;

1509

- It contains "direct" information: CON.D=1;

1510 • The direct information is for itself: CON.DSSID is the SID of the Switch itself.
 1511 Then, the direct switching table entry with the following information is removed:

- 1512 • Uplink (SID, LNID, LCID) = (PKT.SID, PKT.LNID, CON.LCID);
- 1513 • Downlink (SID, LNID, LCID, NAD) = (CON.DCSID, CON.DCLNID, CON.DCLCID, CON.DCNAD).



1514
 1515 **Figure 37 - Release of a direct connection**

1516 **4.3.6.3 Direct connection switching**

1517 As explained in section 4.3.5.2, the normal switching mechanism is intended to be used for forwarding
 1518 communication data between the Base Node and each Service Node. The “direct switching” is a mechanism
 1519 to let two Nodes communicate with each other, switching the packets in a local way, i.e. without passing
 1520 through the Base Node. It is not a different form of packet-switching, but rather an additional feature of the
 1521 general switching process.

1522 The first shared Switch in the paths that go from two Service Nodes to the Base Node will be called the “direct
 1523 Switch” for the connections between the said Nodes. This is the Switch that will have the possibility of
 1524 performing the direct switching to make the two Nodes communicate efficiently. As a special case, every
 1525 Switch is the “direct Switch” between itself and any Node that is lower down in the hierarchy.

1526 The “direct switching table” is a table every Switch should contain in order to perform the direct switching.
 1527 Each entry on this table is a direct connection that must be switched directly. It is represented by the origin

1528 CID and the destination CID of the direct connection. It is not a record of every connection identifier lower
1529 down in its hierarchy, but contains only those that should be directly switched by it. The Destination Node's
1530 ability to receive aggregated packets shall also be included in the "direct switching table" in order to fill the
1531 PKT.NAD field.

1532 **4.3.6.4 Direct switching operation**

1533 If a Switch receives an uplink (HDR.DO=0) MAC frame that is to be switched (see section 4.3.5.2 for the
1534 requirements) and its address is in the direct switching table, then the procedure is as follows:

- 1535 • Change the (SID, LNID, LCID, NAD) by the Downlink part of the entry in the direct switching table.
- 1536 • Queue the packet to be transmitted as a Downlink packet (HDR.DO=1).

1537 **4.3.7 Packet aggregation**

1538 **4.3.7.1 General**

1539 The GPDU may contain one or more packets. The functionality of including multiple packets in a GPDU is
1540 called packet aggregation. Packet aggregation is an optional part of this specification and devices do not need
1541 to implement it for compliance with this specification. It is however suggested that devices should implement
1542 packet aggregation in order to improve MAC efficiency.

1543 To maintain compatibility between devices that implement packet aggregation and ones that do not, there
1544 must be a guarantee that no aggregation takes place for packets whose data transit path from/to the Base
1545 Node crosses (an) intermediate Service Node(s) that do(es) not implement this function. Information about
1546 the aggregation capability of the data transit path is exchanged during the Registration process (4.6.1). A
1547 registering Service Node notifies this capability to the Base Node in the REG.CAP_PA field (1 bit, see Table
1548 20) of its REG_REQ message. It gets feedback from the Base Node on the aggregation capability of the whole
1549 Downlink transit path in the REG.CAP_PA field of the REG_RSP message.

1550 Based on initial information exchanged on Registration, each subsequent data packet in either direction
1551 contains aggregation information in the PKT.NAD field. In the Downlink direction, the Base Node will be
1552 responsible for filling PKT.NAD based on the value it communicated to the destination Service Node in the
1553 REG.CAP_PA field of the REG_RSP message. Likewise, for uplink data, the source Service Node will fill
1554 PKT.NAD based on the REG.CAP_PA field received in the initial REG_RSP from the Base Node. The last Switch
1555 shall use the PKT.NAD field to avoid packet aggregation when forwarding the packet to destination Service
1556 Nodes without packet aggregation capability. Intermediate Switch Nodes should have information about the
1557 aggregation capability in their switching table and shall not aggregate packets when it is known that next
1558 level Switch Node does not support this feature.

1559 Devices that implement packet aggregation shall ensure that the size of the MSDU comprising the aggregates
1560 does not exceed the maximum capacity of the most robust transmission scheme of a PHY burst. The most
1561 robust transmission scheme refers to the most robust combination of modulation scheme, convolutional
1562 coding and repetition coding.

1563 4.3.7.2 Packet aggregation when switching

1564 Switch Nodes maintain information on the packet aggregation capability of all entries in their switching table,
1565 i.e. of all switches that are connected to the Subnetwork through them. This capability information is then
1566 used during traffic switching to/from the connected Switch Nodes.

1567 The packet aggregation capability of a connecting Switch Node is registered at each transit Switch Node at
1568 the time of its promotion by sniffing relevant information in the PRO_ACK message.

- 1569 • If the PKT.SID in a PRO_ACK message is the same as the switching Node, the Node being promoted is
1570 connected directly to the said Switch Node. The aggregation capability of this new Switch Node is
1571 registered as the same as indicated in PKT.NAD of the PRO_ACK packet.
- 1572 • If the PKT.SID in a PRO_ACK message is different from the SID of the switching Node, it implies that
1573 the Node being promoted is indirectly connected to this Switch. The aggregation capability for this
1574 new Switch Node will thus be the same as the aggregation capability registered for its immediate
1575 Switch, i.e. PKT.SID.

1576 Aggregation while switching packets in uplink direction is performed if the Node performing the Switch
1577 knows that its uplink path is capable of handling aggregated packets, based on capability information
1578 exchanged during Registration (REG.CAP_PA field in REG_RSP message).

1579 Downlink packets are aggregated by analyzing the following:

- 1580 • If the PKT.SID is the same as the switching Node, then it is the last switching level and the packet will
1581 arrive at its destination. In this case, the packet may be aggregated if PKT.NAD=0.
- 1582 • If the PKT.SID is different, this is not the last level and another Switch will receive the packet. The
1583 information of whether or not the packet could be aggregated should be extracted from the switching
1584 table.

1585 4.3.8 Security

1586 4.3.8.1 General

1587 The security functionality provides the MAC layer with confidentiality, authentication, integrity and
1588 protection against reply attacks through a secure connection method and a key management policy. All
1589 packets must use the negotiated security profile.

1590 4.3.8.2 Security Profiles

1591 Several security profiles are provided for managing different security needs, which can arise in different
1592 network environments. This version of the specification lists three security profiles and leaves scope for
1593 adding another security profile in future versions.

1594 4.3.8.2.1 Security Profile 0

1595 Communications having Security Profile 0 are based on the transmission of MAC SDUs without encryption.
1596 This profile may be used in application scenarios where either sufficient security is provided by upper
1597 communication layers or where security is not a major requirement for application use-case.

1598 4.3.8.2.2 Security Profile 1 and 2

1599 4.3.8.2.2.1 General

1600 Security Profile 1 and 2 are based on several cryptographic primitives, all based upon AES-128, which provides
1601 secure functionalities for key derivation, key wrapping/unwrapping and authenticated encryption of packets.
1602 These profiles are specified with the aim of fulfilling all security requirements:”

- 1603 • Confidentiality, authenticity and integrity of packets are guaranteed by the use of an authenticated
1604 encryption algorithm.
- 1605 • Authentication is guaranteed by the fact that each Node has its own unique key known only by the
1606 Node itself and the Base Node.
- 1607 • Replay Attacks are prevented through the use of a message counter of 4 bytes.

1608 **Note:**

1609 The scope of the Security Profile does not address any implementation specific security requirements such
1610 as protection against side channel attacks (timing attacks, power attacks, electromagnetic attacks, fault
1611 attacks, etc...). The implementer of the security profile needs to assure the cryptographic functionality is
1612 adequately protected.

- 1613 • The implementer might consider counter measures depending on the environment PRIME is
1614 used. This could include the implementation of an AES algorithm with mitigation for non-invasive
1615 attacks (e.g. power analysis or electromagnetic side channel attacks). Additional tamper
1616 protection and hardening mechanisms are specified in FIPS 140-3 levels 3 and 4.

1617 4.3.8.2.2.2 Authenticated Encryption

1618 The cryptographic algorithms used in this specification are all based on AES, as specified in [16]. The
1619 specification describes the algorithm with three possible key sizes. PRIME uses a key length of 128 bit. A key
1620 length of 128 bit represents a good level of security for preserving privacy up to 2030 and beyond, as specified
1621 in SP800-57 [17], page 66, table 4.

1622 AES is used in CCM mode, as specified in [25]. It is a dual-pass authenticated encryption mode. In the context
1623 of this security profile it is used accordingly to the following settings (using the same notations of [25]):

- 1624 • *n*: the octet length of the nonce is set to 13. This allows for a maximum message size of 65535
1625 bytes.
- 1626 • *q*: the octet length of the binary representation of the octet length of the payload is set to 2.
- 1627 • *t*: the octet length of the MAC is set to 6. Therefore *Tlen*, the MAC bit size, is set to 48.

1628 4.3.8.2.2.2.1 Key update frequency

1629 Security profiles set the value of the AES-CCM authentication tag (*Tlen*) to 48 bits. The maximum time limit
1630 between two re-keying events for WK and SWK is the value contained in the PIB *MACUpdateKeysTime*. This
1631 PIB's maximum allowed value shall be the one defined in Table 13 according to the available number of
1632 channels.

Table 13 - Values of *MACUpdateKeysTime* for different number of channels

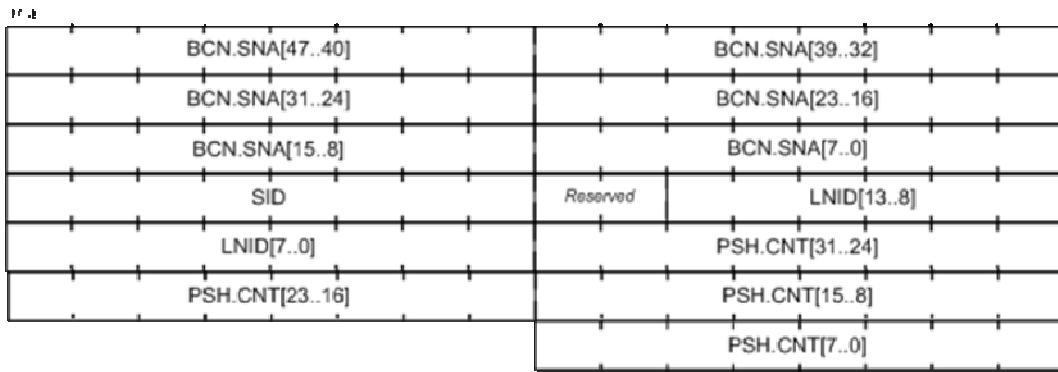
Available number of channels	Life time in days
1 x 64Kb/s channel	49 days
2 x 64Kb/s channel	24 days
3 x 64Kb/s channel	16 days
4 x 64Kb/s channel	12 days
5 x 64Kb/s channel	10 days
6 x 64Kb/s channel	8 days
7 x 64Kb/s channel	7 days
8 x 64Kb/s channel	6 days

1634 **4.3.8.2.2.2.2 Nonce creation**

1635 The nonce is a value used by AES-128-CCM and is required to be unique for each different message that is
 1636 processed under the same key. In order to maintain this property and to have protections against replay
 1637 attacks, each Service Node needs to have a 32-bit message counter.

1638 The 13-byte nonce, for each message, is shown in Figure 38 and is composed by the concatenation of the
 1639 following entities:

- 1640 • 48-bit Subnetwork Address (found in BCN.SNA)
- 1641 • 8-bit SID address, identifying the Switch Node of the Service Node which generated the packet
- 1642 • 2-bit set to 0 for this version of the specification. Reserved for future use.
- 1643 • 14-bit LNID address, identifying the Service Node that generated the packet. The pair SID and
 1644 LNID should provide a unique address within the subnetwork.
- 1645 • 32-bit Message Counter, number of messages sent by the Service Node which originated the
 1646 message



1647

1648

Figure 38 - Nonce structure

1649 The nonce SID and LNID entities are derived from the senders SID and LNID as shown in Table 14.

1650

Table 14 -Nonce SID/LNID Derivation

Packet Type / Affected Keys	Packet Direction	Nonce Entity	Nonce Entity Value
Unicast Key: WK or SWK not re-encoded	Downlink	SID	0
		LNID	0
	Uplink	SID	PKT.SID
		LNID	PKT.LNID
Switch Re-encoded Key: SWK *	Downlink/	SID	Switch SSID
	Uplink	LNID	0
Broadcast/Multicast Key: SWK	Downlink	SID	PKT.SID
		LNID	0
	Uplink	SID	PKT.SID
		LNID	PSH.LNID

1651 * Switch re-encoded packets are all ALV messages sent by switches as well as all downlink messages encrypted
1652 with SWK. A switch shall know the immediate switch for each SID below it to be able to derive the nonce.

1653 **4.3.8.2.2.2.1 Message counter (PSH: CNT)**

1654 In order to avoid repetition attacks, in case the message counter of a message is not correctly validated
1655 according to this chapter, the message shall be discarded.

1656 In the case of messages protected with WK, each node has a message counter starting from zero,
1657 incrementing after each protected message sent. This counter shall be reset after updating to a newer key.

1658 In the case of messages protected with SWK, the counter used for nonce creation shall act following a
1659 distributed algorithm. This counter shall also be reset after updating to a newer key.

1660 Upon reception of downlink message, every node shall validate that the counter is bigger than the last
1661 downlink message received. If the node is an intermediate Switch Node, once validated, the message shall
1662 be re-encoded before switching it down.

1663 Upon reception of uplink message, a Switch Node shall not perform a validation of the CNT except for the
1664 cases when the message is processed by them. These are the use cases when it shall validate the message:

- 1665
- ALV messages
 - Data uplink messages switched down by a Direct Switch (does not apply to the intermediate Switch
1667 Nodes, only the Direct Switch itself)

1668 A node shall maintain a transmission counter with the next value to be used in CNT field. This value shall be
1669 used when the node encodes and re-encodes a message. After using it, the node shall increase its value by
1670 1.

1671 When the node receives a counter bigger than the one stored for transmission, the transmission counter
1672 shall be updated with the received counter value + 1. This applies to both downlink and uplink messages. As
1673 a guideline, note that this forces any response message to have a counter bigger than the related request
1674 message, easing validation requirement and corner cases.

1675 In security profile 1, direct switches are allowed to ignore validation of the first uplink data message from
1676 each peer, after that first message, the switch shall perform a validation of all uplink messages. This is to
1677 simplify implementation of direct switches on both memory usage and complexity.

1678 Base Node shall validate all the uplink counters, considering all the rules described previously. In order to
1679 perform this, it shall keep track of individual counters for each node. The algorithm to do so is left up to the
1680 manufacturer of the Base Node.

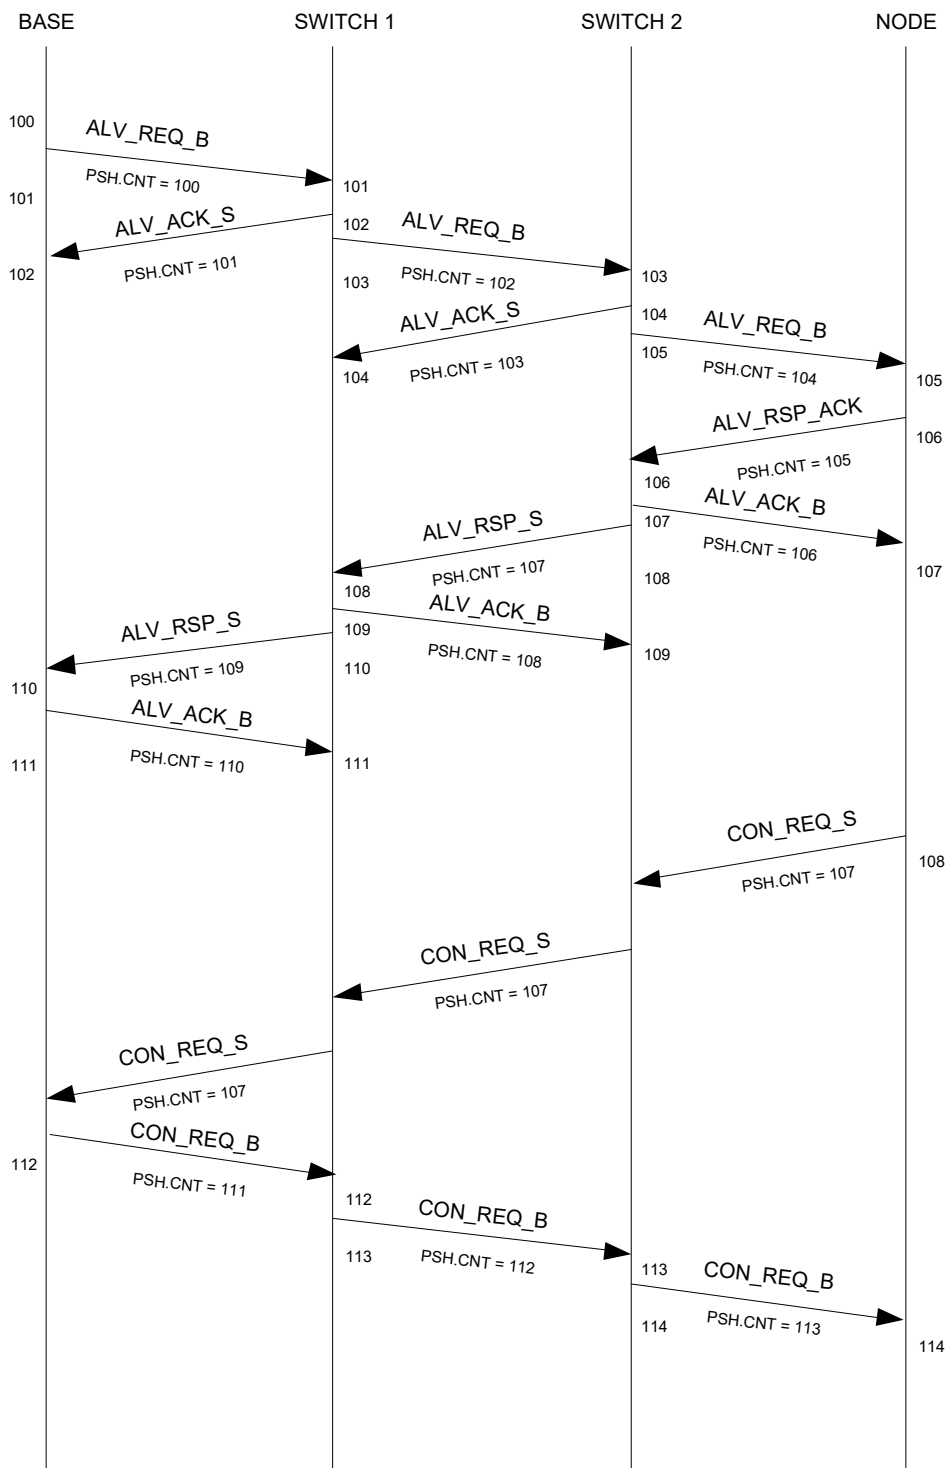


Figure 39 - Counter handling in ALV and request/response messages

1681

1682

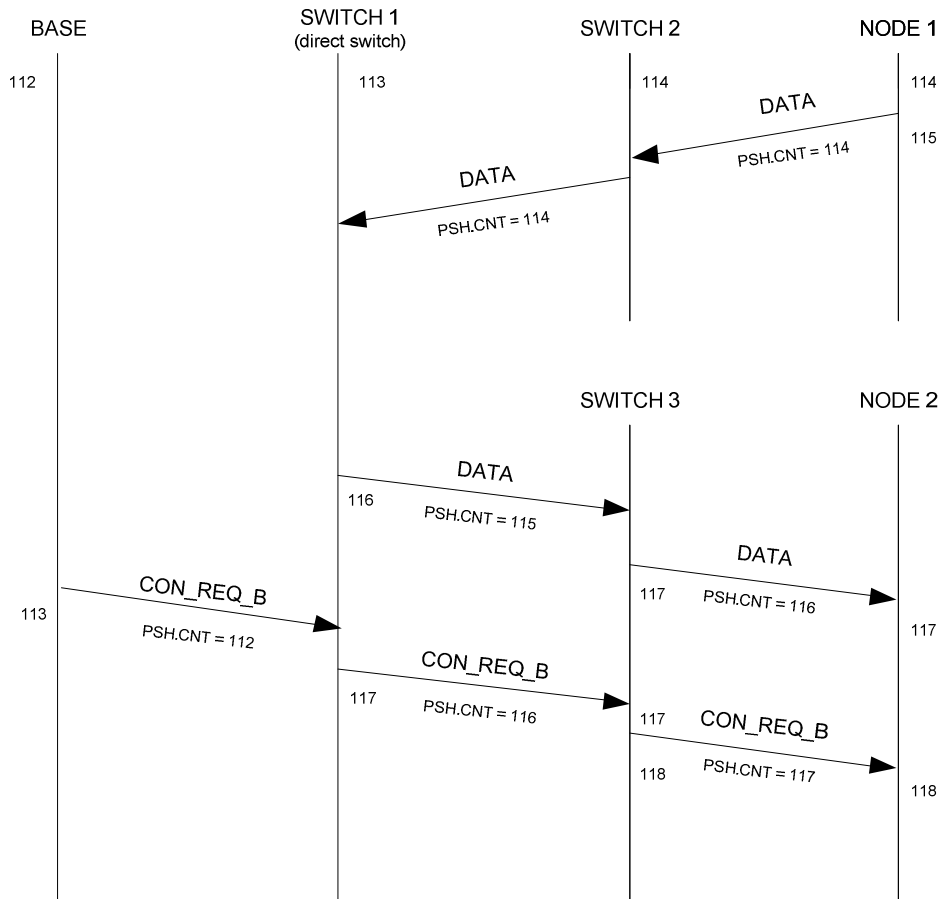



Figure 40 - Counter handling in direct connections and request messages

4.3.8.2.2.2.3 Creation of Challenge nonce for REG PDU-s

Each REG message relies on a 64 bit random number and the Subnetwork Address as a challenge. This challenge shall be used as the nonce for the authentication of the REG_REQ, REG_RSP and REG_REJ messages. The usage sequence is described in detail in 4.6.1.2.

The nonce, for each message, is shown in Figure 41, and is composed by the concatenation of the following entities:

- 40 least significant bits of the Subnetwork Address (found in BCN.SNA).
- For REG_REQ, Service Node shall generate a 64-bit random number and shall encode that value in the PSH.CNT and REG.CNT fields.
- For REG_RSP and REG_REJ, REG.CNT shall be the copy of the same field coming from REG_REQ and PSH.CNT shall be PSH.CNT coming from REG_REQ incremented by 1, overflowing if necessary.



1697

1698

Figure 41 - REG Nonce structure

1699 **4.3.8.2.2.3 Key Derivation Algorithm**

1700 The method for key derivation is KDF in counter mode as specified in [23] using AES-CMAC [24] with key size
 1701 of 128 as underlying PRF. This KDF requires 5 values as input:

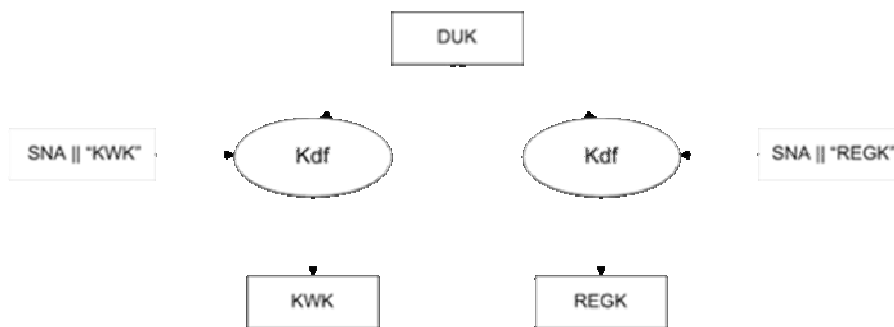
- 1702 • K_i which is the master key used to derive the output key K_o
- 1703 • *Label* which is a string, fixed for the purpose of this security profile at “PRIME_MAC”
- 1704 • *Context*, which is a string assuming different values accordingly to the purpose of the output key,
 1705 which will be described in section 4.3.8.3.2
- 1706 • L which is the size of the output key, which for the purpose of this security profile is fixed to 128.
- 1707 • r which is an integer indicating the lengths of the binary representation of the counter and of L ,
 1708 which is fixed in this security profile to 32

1709 **4.3.8.2.2.4 Key Derivation Hierarchy**

1710 Figure 42 outlines the Key Derivation hierarchy and the process to derive the Key Wrapping Key (KWK) and
 1711 the Registration Key (REGK).

1712 The KWK is used to wrap the individual Working Key (WK) and the Subnetwork Working Key (SWK) when sent
 1713 down from the Base Node to the Terminal Node, while the REGK is used for authentication in the registration
 1714 process to authenticate both, BN and TN.

1715 The random number generator used for WK and SWK should be compliant with [27].



1716

1717

Figure 42 - Key derivation hierarchy

1718 **4.3.8.2.2.5 Key Wrapping Algorithm**

1719 The method for wrapping and unwrapping keys is referred to AES-128-KW, it is described as KW in [26], and
1720 uses AES-128 as underlying cipher. It is used to transmit keys in an encrypted form. In this security profile all
1721 keys are of 128 bit, which means that wrapped keys are 192 bits.

1722 **4.3.8.2.3 Encryption/Authentication by PDU Types**

1723 The following table shows which PDU-s are authenticated (A) and/or encrypted (E) on each of the security
1724 profiles. This table shows packet types with their names as presented in section 4.3.8.6. The packet
1725 nomination follows the following rules: if the packet is a generic name (e.g. REG), the profile will apply for all
1726 the subpacket types not listed in the table (e.g. REG_ACK).

1727 **Table 15 - Encryption/Authentication by PDU Types**

PDU Type	Profile 1	Profile 2
REG_REQ, REG_RSP	REGK (A)	REGK (A)
REG_REJ	Plain	REGK (A)
Unicast DATA	WK (AE)*	WK (AE)*
SEC	WK(AE)	WK(AE)
Multicast DATA, Broadcast DATA, Direct connection DATA	SWK (AE)*	SWK (AE)*
PRO, MUL, CFP, CON, FRA, ALV, PRO_ACK, MUL_JOIN_B, MUL_LEAVE_S, PRO_DEM_S, PRO_DEM_B, REG_UNR_B, REG_UNR_S	Plain	SWK (AE)
REG_ACK	Plain	WK(A)

1728 The rows highlighted with an asterisk (*) can be optionally sent not encrypted as described in section 4.3.8.6.

1729 **4.3.8.3 Negotiation of the Security Profile**

1730 **4.3.8.3.1 General**

1731 All MAC data, including signaling PDUs (all MAC control packets defined in section 0) use the same security
1732 profile. This profile is negotiated during the device Registration. In the REG_REQ message the Terminal
1733 indicates a security profile it is able to support in the field REG.SPC. The Base Node may accept this security
1734 profile and so accept the Registration, sending back a REG_RSP with the same REG.SPC value. The Base Node
1735 may also accept the Registration, however it sets REG.SPC to 0, 1 or 2 indicating that security profile 0, 1 or
1736 2 is to be used. Alternatively, the Base Node may reject the Registration if the Terminal does not provide an
1737 acceptable security profile.

1738 It is recommended that the Terminal first attempts to register using the highest security profile it supports.
1739 In case the Base Node replies with a different value for REG.SPC, corresponding to a profile with lower
1740 security, the Terminal could refuse the registration by not sending the REG_ACK. The policy used by the

1741 Terminal to refuse a registration with a lower than expected security profile is out of the scope of this
1742 specification.

1743 4.3.8.3.2 Key Types and Key Hierarchy

1744 The key hierarchy of Security Profile 1 and 2 is based on three assumptions:

- 1745 1. There is a 128 bit unique key on each service node called Device Unique Key (DUK). How this key is
1746 generated, provided to service nodes, is out of the scope of this specification. The DUK is managed
1747 by macSecDUK (refer to section 6.2.3.6.)
- 1748 2. The Base Node must have knowledge of a Service Node's DUK by only knowing its EUI-48.
- 1749 3. As specified by [REF TO NIST SP800-57Part1] "In general, a single key should be used for only one
1750 purpose".

1751 The keys and their respective usage are:

1752 **Device Unique Key (DUK):** DUK is used only for key derivation purposes, using the KDF described in section
1753 4.3.8.2.2.3. It has the requirement to be unique for each device. It is used to generate KWK and REGK.

1754 **Key Wrapping Key (KWK):** This key is derived from DUK using the concatenation of the Subnetwork Address
1755 (SNA) and the string "KWK" as *Context*. It is used to unwrap the keys received from the Base Node.

1756 **REG Key (REGK):** This key is derived from DUK using the concatenation of the Subnetwork Address (SNA) and
1757 the string "REGK" as *Context*. It is used to protect, through AES-128-CCM, some of the REG control messages,
1758 specifically it is used for: REG_REQ, REG_RSP, REG_REJ only when REG.R=0. The reason is that there hasn't
1759 been any communication with the Base Node yet, so no other shared keys have been established.

1760 **Working Key (WK):** This key is used to encrypt all the unicast data that is transmitted from the Base Node to
1761 a Service Node and vice versa. Each registered Service Node would have a unique WK that is known only to
1762 the Base Node and itself. The WK is randomly generated by the Base Node, wrapped through AES-128-KW
1763 and transmitted by the Base Node in REG_RSP and SEC messages.

1764 **Subnetwork Working Key (SWK):** The SWK is shared by the entire Subnetwork. The SWK is randomly
1765 generated by the Base Node, wrapped through AES-128-KW and transmitted by the Base Node in REG_RSP
1766 and SEC messages.

1767 The WK and the SWK have a limited validity time related to the random sequence generation period. The
1768 random sequence is regenerated and distributed by the Base Node at least every *MACUpdateKeysTime*
1769 seconds through the SEC control packet. If a device does not receive a new SEC message within
1770 *MACUpdateKeysTime* it shall move back from its present functional state to a *Disconnected* functional state.

1771 The key hierarchy has been designed to ensure security of the required MAC keys, to follow NIST
1772 specifications and to be as simple as possible.

1773 4.3.8.4 Key Distribution and Management

1774 The Security Profile for data traffic is negotiated when a device is registered. The REG control packet contains
1775 specific fields to indicate the Security Profile for respective devices. All connections to/from the device would

1776 be required to follow the Security Profile negotiated at the time of Registration. There cannot be a difference
1777 in Security Profile across multiple connections involving the same device. The only exception to this would
1778 be the Base Node.

1779 All keys are never transmitted in non-encrypted form over the physical channel. The SEC unicast messages
1780 transmitted by the Base Node at regular intervals contain random keys for both unicast and non-unicast
1781 traffic.

1782 When a device initially registers on a Subnetwork, the REG response from the Base Node contains the
1783 wrapped SWK and WK. If the SN cannot unwrap the keys successfully, it shall discard the REG response and
1784 continue trying to register.

1785 The process of updating WK is as follows:

- 1786 • The SN shall start using the new WK immediately for transmission.
- 1787 • The BN shall use the old WK for transmission until receiving the SEC response.
- 1788 • The SN and BN shall be able to receive messages encrypted with both new and old WK until SEC
1789 exchange completes or Control packet retransmission (see 4.4.2.6.2) completes.
- 1790 • If the SN cannot unwrap the new WK successfully, it shall indicate that it could not update WK and
1791 shall continue using the old WK. In such case, the BN shall retransmit the SEC message.

1792 Upon reception of a new SWK that is successfully unwrapped, the node shall maintain the old SWK and use
1793 it to decrypt and encrypt the appropriate messages until:

- 1794 1. a message is received encrypted with the new SWK.
- 1795 2. the expiration of a 3-hour timer.

1796 The timer provides a method for ensuring the old SWK will not be used indefinitely.

1797 If the SN cannot unwrap the new SWK successfully, it shall indicate that it could not update SWK and shall
1798 continue using the old SWK. In such case, the BN shall retransmit the SEC message.

1799 It is recommended that a Base Node register a Terminal Node with the old SWK and immediately perform a
1800 SEC procedure, with the registering Terminal Node, if the Terminal Node registers during an in progress SWK
1801 SEC procedure.

1802 The Base Node shall maintain the old SWK for duration not to exceed 3-hours from the beginning of the SWK
1803 SEC procedure. This limit constrains the SWK SEC procedure duration to 3-hours. The Base Node can stop
1804 accepting the old SWK from any node before that duration, e.g. if it is certain that a Service Node has received
1805 a packet with the new SWK.

1806 **4.3.8.5 Encryption and Authentication**

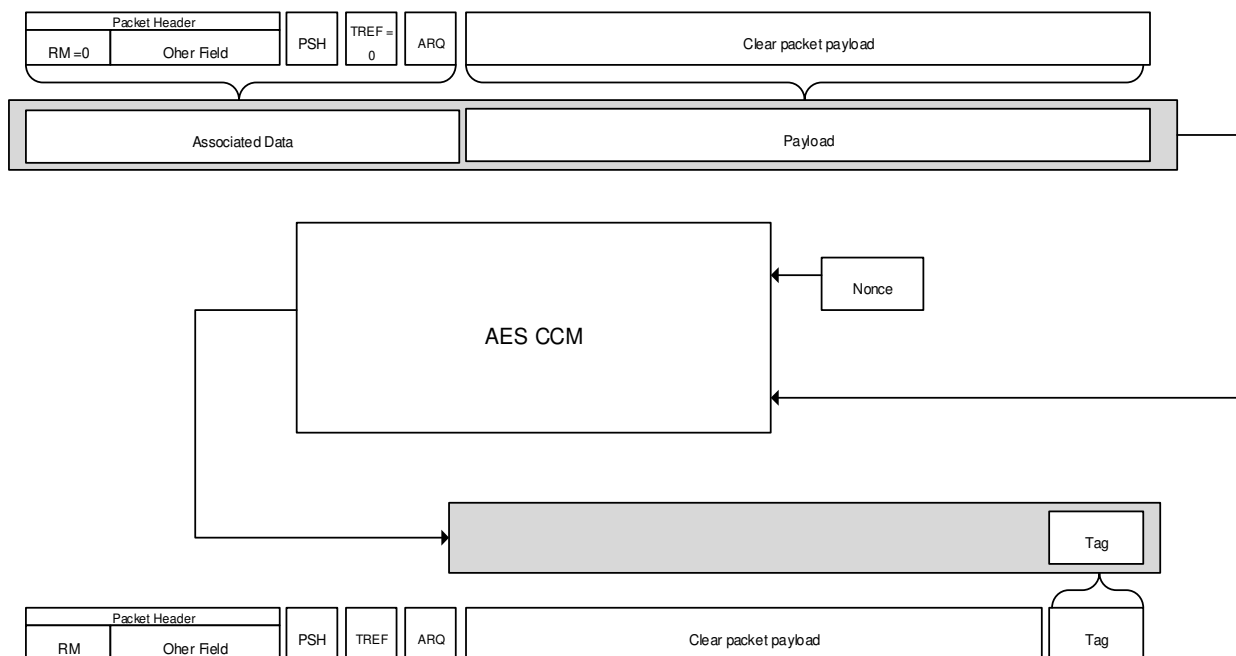
1807 **4.3.8.5.1 Security Profile 0**

1808 Not Applicable.

1809 **4.3.8.5.2 Security Profile 1 and 2**

1810 Security Profiles 1 and 2 make use of AES-CCM for packet protection but there are three different cases,
 1811 accordingly to section 4.3.8.2.3:

- 1812 • Plain: in the case the packet is not processed by AES-CCM and there is no Tag
- 1813 • Authentication Only: in this case the packet header where the PKT.RM is masked with 0, PSH, TREF
 1814 (if exists) where the values are masked with 0, ARQ (if exists) and the payload should be processed by
 1815 AES-CCM as associated data. This situation is depicted in Figure 43.
- 1816 • Authentication and Encryption: in this case the packet header where the PKT.RM is masked with 0,
 1817 PSH, TREF (if exists) where the values are masked with 0 and ARQ (if exists) should be processed as
 1818 associated data, thus only being authenticated, while the payload should be processed as payload,
 1819 thus authenticated and encrypted. This situation is depicted in Figure 44.



1820
 1821

Figure 43 - Security profile 1 and 2 encryption algorithm (authentication only)

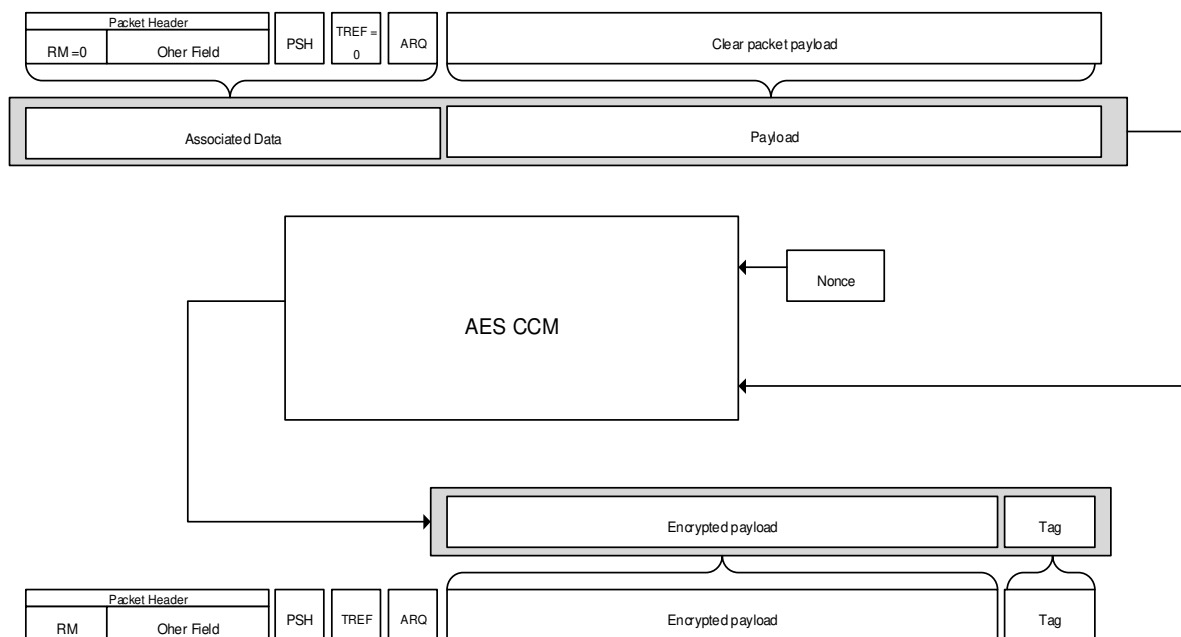


Figure 44 - Security profile 1 and 2 encryption algorithm (authentication and encryption)

4.3.8.6 Unicast and Multicast connection security negotiation

There are some use cases in which it is not desirable to authenticate and encrypt data packets. One use case of this is the firmware upgrade for images that are already signed and potentially encrypted.

For such cases both the unicast and multicast connection establishment procedures have a feature that allows a negotiation for enabling or disabling the security on DATA packets. This procedure only applies to profiles 1 and 2, because profile 0 does not allow encryption.

Each side shall indicate in the CON_REQ and MUL_JOIN packets if the DATA packets for that connection shall be authenticated and encrypted or not.

- In case both sides indicate that the DATA packets shall not be securitized, then the packets shall be sent without being authenticated or encrypted.
- If at least one of the sides indicates that the packets shall be securitized, then the packets shall be authenticated and encrypted.

Regarding broadcast connections, no negotiation takes place. Broadcast connections shall always be authenticated and encrypted.

4.3.8.7 Unicast and Multicast connection security negotiation

After initial negotiation, any multicast connection sending plain data can be upgraded to be authenticated and encrypted at any point. The usual use case is that a new node has joined the multicast group negotiating the connection to be authenticated and encrypted.

The upgrade mechanism is to start sending multicast data authenticated and encrypted, without any additional control message. Any node receiving multicast data that is authenticated and

1844 encrypted shall assume the connection has been upgraded to be authenticated and encrypted
 1845 from the reception of that data.

1846 There is no procedure for a multicast connection to be downgraded to send plain data.

1847 4.4 MAC PDU format

1848 4.4.1 General


1849 There are different types of MAC PDUs for different purposes.

1850 4.4.2 Generic MAC PDU

1851 4.4.2.1 General

1852 Most Subnetwork traffic comprises Generic MAC PDUs (GPDU). GPDUs are used for all data traffic and most
 1853 control traffic. All MAC control packets are transmitted as GPDUs.


1854 GPDU composition is shown in Figure 45. It is composed of a Generic MAC Header followed by one or more
 1855 MAC packets and 32 bit CRC appended at the end.



1856
 1857 **Figure 45 - Generic MAC PDU format**

1858 4.4.2.2 Generic MAC Header

1859 The Generic MAC Header format is represented in Figure 46 and Table 16. The size of the Generic MAC
 1860 Header is 3 bytes. Table 16 enumerates each field of a Generic MAC Header.



1861
 1862 **Figure 46 - Generic MAC header**


1863 **Table 16 - Generic MAC header fields**

Name	Length	Description
Unused	2 bits	Unused bits those are always 0; included for alignment with MAC_H field in PPDU header (Section 3.4.3).
HDR.HT	2 bits	Header Type. HDR.HT = 0 for GPDU

Name	Length	Description
<i>Reserved</i>	5 bits	Always 0 for this version of the specification. Reserved for future use.
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 $M(x)$.

1864 **4.4.2.3 Packet structure**

1865 A packet is comprised of a Packet Header and Packet Payload. Figure 36 shows the structure.

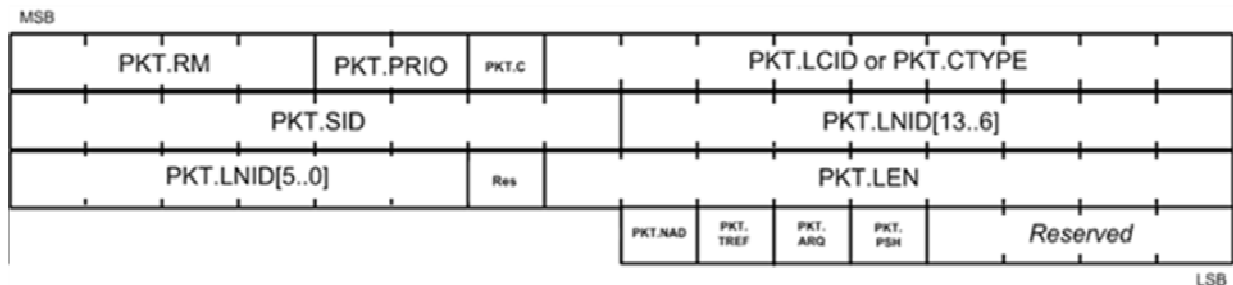


1866

1867

Figure 47 - Packet structure

1868 Packet header is 7 bytes in length and its composition is shown in Figure 48. Table 17 enumerates the
1869 description of each field.

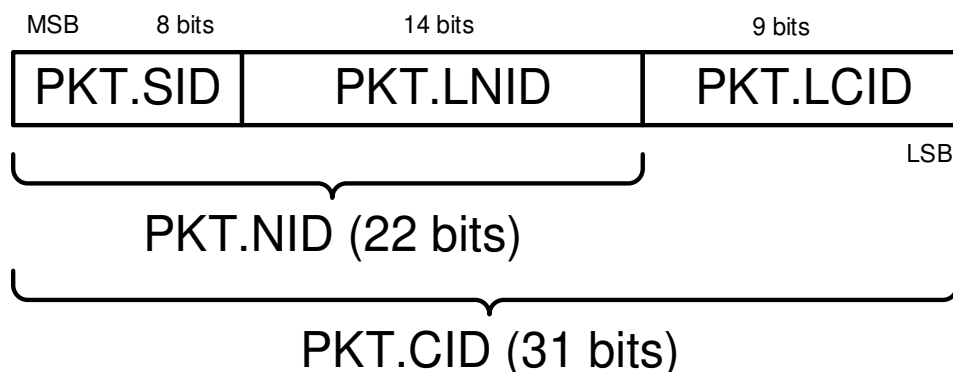


1870

1871

Figure 48 - Packet Header

1872 To simplify, the text contains references to the PKT.NID fields as the composition of the PKT.SID and PKT.LNID.
 1873 The field PKT.CID is also described as the composition of the PKT.NID and the PKT.LCID. The composition of
 1874 these fields is described in Figure 49.



1875
 1876
 1877

Figure 49 - PKT.CID structure

Table 17 - Packet header fields

Name	Length	Description
PKT.RM	4 bits	Weakest modulation this node can decode from the receiving peer: <ul style="list-style-type: none"> • 0 – DBPSK • 1 – DQPSK • 2 – D8PSK • 3 – Not used • 4 – DBPSK + Convolutional Code • 5 – DQPSK + Convolutional Code • 6 – D8PSK + Convolutional Code • 7-11 – Not used • 12 – Robust DBPSK • 13 – Robust DQPSK • 14 – Not used • 15 – Outdated information
PKT.PRIO	2 bits	Indicates 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 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.

Name	Length	Description
PKT.RM	4 bits	Weakest modulation this node can decode from the receiving peer: <ul style="list-style-type: none"> • 0 – DBPSK • 1 – DQPSK • 2 – D8PSK • 3 – Not used • 4 – DBPSK + Convolutional Code • 5 – DQPSK + Convolutional Code • 6 – D8PSK + Convolutional Code • 7-11 – Not used • 12 – Robust DBPSK • 13 – Robust DQPSK • 14 – Not used • 15 – Outdated information
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.
Reserved	1bit	Always 0 for this version of the specification. Reserved for future use.
PKT.LEN	9 bits	Length of the packet excluding the packet header and the authentication tag (if present). It is the sum of the lengths of the payload and the subheaders (if any).
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.TREF	1 bit	TREF subheader presence: <ul style="list-style-type: none"> • If PKT.TREF=0 the packet doesn't include a TREF subheader. • If PKT.TREF=1 the packet includes a TREF subheader.
PKT.ARQ	1 bit	ARQ subheader presence: <ul style="list-style-type: none"> • If PKT.ARQ=0 the packet doesn't include an ARQ subheader. • If PKT.ARQ=1 the packet includes an ARQ subheader.
PKT.PSH	1 bit	Packet security subheader presence: <ul style="list-style-type: none"> • If PKT.PSH=0 the packet doesn't include a security subheader. • If PKT.PSH=1 the packet includes a security subheader.
Reserved	4 bits	Always 0 for this version of the specification. Reserved for future use.

1878 The "ARQ subheader", "TREF subheader" and "security subheader" are optional. Their presence depends on
1879 the PKT.ARQ, PKT.TREF and PKT.PSH flags. The description of the ARQ subheader will be done in the section

1880 4.7.3.2 and the description of the TREF subheader will be done in section 4.8. MAC Control packets shall not
 1881 include a TREF or ARQ subheader.

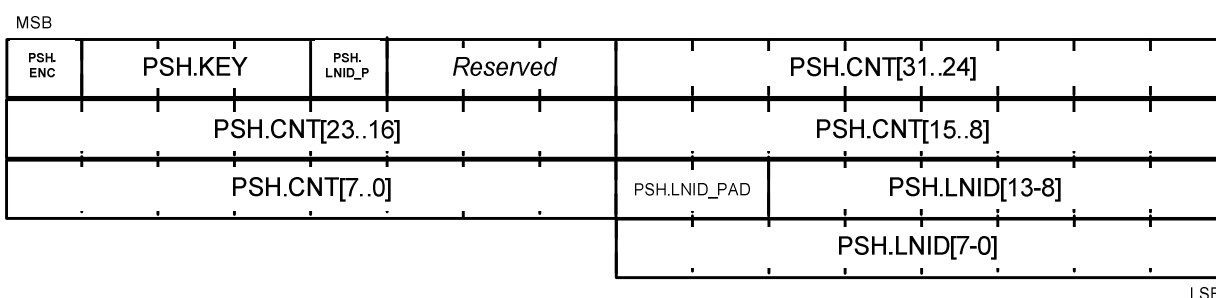
1882 **4.4.2.4 CRC**

1883 The CRC is the last field of the GPDU. It is 32 bits long. It is used to detect transmission errors. The CRC shall
 1884 cover the concatenation of the SNA with the GPDU except for the CRC field itself.

1885 The input polynomial $M(x)$ is formed as a polynomial whose coefficients are bits of the data being checked
 1886 (the first bit to check is the highest order coefficient and the last bit to check is the coefficient of order zero).
 1887 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
 1888 remainder $R(x)$ is calculated as the remainder from the division of $M(x)\cdot x^{32}$ by $G(x)$. The coefficients of the
 1889 remainder shall then be the resulting CRC.

1890 **4.4.2.5 Security header**

1891 For the security profiles 1 and 2, the security subheader contains the needed information to authenticate
 1892 and/or encrypt the packet.



1893

1894

Figure 50 - Security subheader

1895 The description of the fields is described in the following table.

1896

Table 18 - Security subheader fields

Name	Length	Description
PSH.ENC	1 bit	Flag to determine if the packet is encrypted: 0 – The packet is Authenticated 1 – The packet is Authenticated and Encrypted
PSH.KEY	3 bits	Key used for the encoding of this packet: 0 – WK 1 – SWK 2 – REG 3-8 – Reserved for future used.

Name	Length	Description
PSH.LNID_P	1 bit	Flag to determine if the PSH.LNID is present (counting the reserved bits leading it). 0 – If PSH.LNID is not present 1 – If PSH.LNID is present
Reserved	3 bits	Always 0 for this version of the specification. Reserved for future use.
PSH.CNT	32 bits	Counter to be used in the nonce composition. * For replay protection, receiving node needs to discard packets that do not follow the rules described in 0.
PSH.LNID_PA D	2 bits	Always 0 for this version of the specification. Reserved for future use. Only present if PSH.LNID_P field set to one.
PSH.LNID	14 bits	Transmitter LNID field to create the nonce when it cannot be derived from the packet. The exception being REG packets. Only present if PSH.LNID_P field set to one.

1897 When the security header is present, a 48-bit authentication tag is appended to the packet. The
1898 authentication tag is the output of the AES-CCM operation (see Figure 44).

1899 4.4.2.6 MAC control packets

1900 4.4.2.6.1 General

1901 MAC control packets enable a Service Node to communicate control information with their Switch Node,
1902 Base Node and vice versa. A control packet is transmitted as a GPDU and is identified with PKT.C bit set to 1
1903 (See section 4.4.2 for more information about the fields of the packets).

1904 There are several types of control messages. Each control message type is identified by the field PKT.CTYPE.
1905 Table 19 lists the types of control messages. The packet payload (see section 4.4.2.3) shall contain the
1906 information carried by the control packets. This information differs depending on the packet type.

1907 **Table 19 - MAC control packet types**

Type (PKT.CTYPE)	Packet name	Packet description
1	REG	Registration management
2	CON	Connection management
3	PRO	Promotion management
5	FRA	Frame structure change
6	CFP	Contention-Free Period request
7	ALV	Keep-Alive

Type (PKT.CTYPE)	Packet name	Packet description
8	MUL	Multicast Management
10	SEC	Security information

1908 **4.4.2.6.2 Control packet retransmission**

1909 For recovery from lost control messages, a retransmit scheme is defined. MAC control transactions
 1910 comprising of exchange of more than one control packet may follow the retransmission mechanism
 1911 described in this section.

1912 The retransmission scheme shall be applied to the following packets when they require a response:

- 1913 • CON_REQ_S, CON_REQ_B;
- 1914 • CON_CLS_S, CON_CLS_B;
- 1915 • REG_RSP;
- 1916 • PRO_REQ_B;
- 1917 • MUL_JOIN_S, MUL_JOIN_B;
- 1918 • MUL_LEAVE_S, MUL_LEAVE_B;
- 1919 • MUL_SW_LEAVE_B
- 1920 • SEC_REQ

1921 Devices involved in a MAC control transaction using retransmission mechanism shall maintain a retransmit
 1922 timer and a message fail timer.

1923 At the requester of a control message transaction:

- 1924 • When the one of the above messages in a transaction is transmitted, the retransmit timer is
 1925 started with value greater or equal to macMinCtlReTxTimer and the control message fail timer is
 1926 started with value macCtrlMsgFailTime.

1927 If a response message is received the retransmit timer and control message fail timer are stopped and the
 1928 transaction is considered complete. Note that it is possible to receive further response messages. These
 1929 would be messages that encountered network delays.

- 1930 • If the retransmit timer expires the control message is retransmitted and the retransmit timer is
 1931 re-started with value greater or equal to macMinCtlReTxTimer (value can be different from the
 1932 previous one).
- 1933 • If the control message fail timer expires, failure result corresponding to respective MAC-SAP
 1934 should be returned to the calling entity. Implementations may also choose to inform their local
 1935 management entity of such failure. If the retransmission is done by the Service Node, the device
 1936 shall return to the *Disconnected* functional state

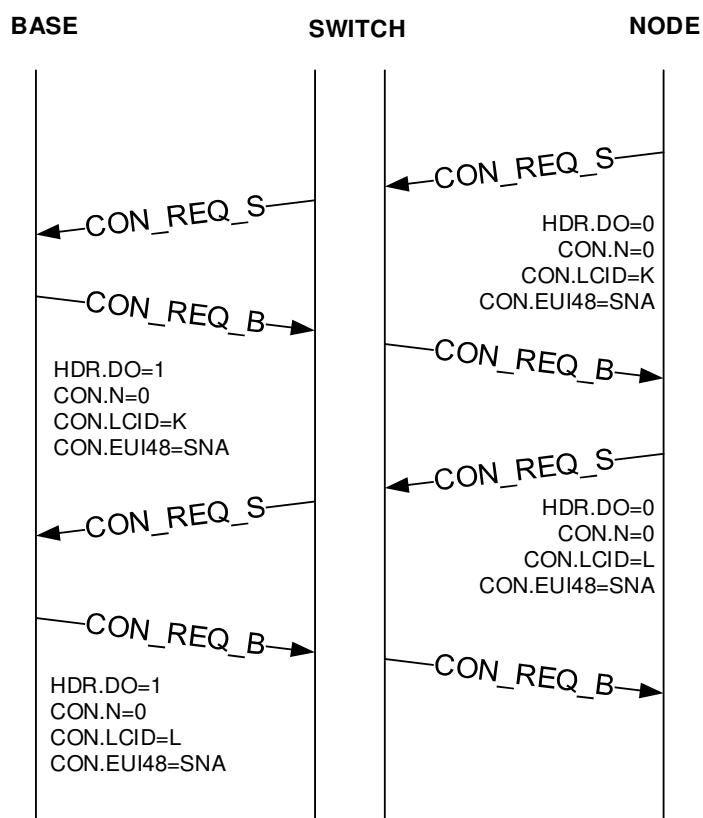
1937 At the responder of a control message transaction:

- 1938 • The receiver of a message must determine itself if this message is a retransmit. If so, no local
 1939 action is needed other than sending a reply to the response.

1940 If the received message is not a retransmit, the message shall be processed and a response returned to the
1941 sender.

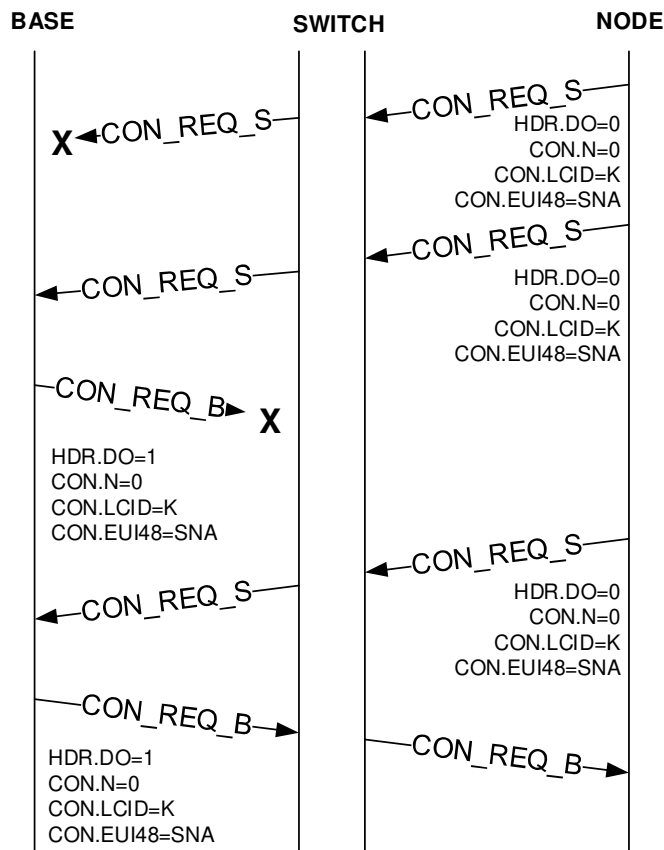
- 1942
- For transactions which use three messages in the transaction, e.g. promotion as shown in 4.6.3,
1943 the responder shall perform retransmits in exactly the same way as the requester. This ensures
1944 that if the third message in the transaction is lost, the message shall be retried and the
1945 transaction completed.

1946 The following message sequence charts show some examples of retransmission. Figure 51 shows two
1947 successful transactions without requiring retransmits.



1948
1949 **Figure 51 - Two transactions without requiring retransmits**

1950 Figure 52 shows a more complex example, where messages are lost in both directions causing multiple
1951 retransmits before the transaction completes.



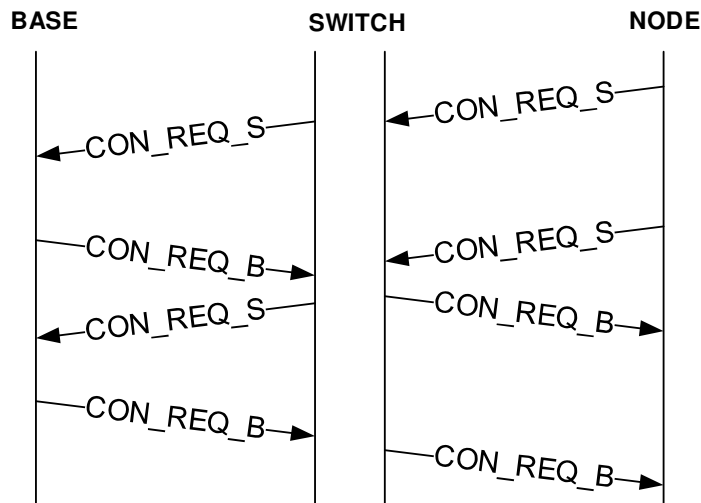
1952

1953

Figure 52 - Transaction with packet loss requiring retransmits

1954

Figure 53 shows the case of a delayed response causing duplication at the initiator of the control transaction.



1955

1956

Figure 53 - Duplicate packet detection and elimination

1957

4.4.2.6.3 REG control packet (PKT.CTYPE=1)

1958

This control packet is used to negotiate the Registration process. The description of data fields of this control

1959

packet is described in Table 20 and Figure 54. The meaning of the packets differs depending on the direction

1960 of the packet. This packet interpretation is explained in Table 20. These packets are used during the
 1961 registration and unregistration processes, as explained in 4.6.1 and 4.6.2.

1962 **Table 20 - REG control packet fields**

Name	Length	Description
REG.N	1 bit	Negative REG.N=1 for the negative register; REG.N=0 for the positive register. (see Table 19)
REG.R	1 bit	Roaming 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. It shall be set by SN in REG_REQ.
REG.SPC	2 bits	Security Profile Capability for Data PDUs: REG.SPC=0 No encryption capability; REG.SPC=1 Security profile 1 capable device; REG.SPC=2 Security profile 2 capable device; REG.SPC=3 Security profile 3 capable device (not yet specified).
REG.CAP_R	1 bit	Robust mode Capable 1 if the device is able to transmit/receive robust mode frames; 0 if the device is not. It shall be set by SN in REG_REQ

Name	Length	Description
REG.CAP_BC	1 bit	<p>Backwards Compatible with 1.3.6</p> <p>1 if the device can operate in backwards compatible mode with 1.3.6 PRIME; 0 if the device is not.</p> <p>It shall be set by SN in REG_REQ and by BN in REG_RSP.</p>
REG.CAP_SW	1 bit	<p>Switch Capable</p> <p>1 if the device is able to behave as a Switch Node; 0 if the device is not.</p> <p>It shall be set by SN in REG_REQ.</p>
REG.CAP_PA	1 bit	<p>Packet Aggregation Capability</p> <p>1 if the SN device has packet aggregation capability; if the BN device has packet aggregation capability together with all the switches along the downlink path to the registration requesting SN and the requesting SN itself; 0 otherwise.</p> <p>It shall be set by SN in REG_REQ and by BN in REG_RSP.</p>
REG.CAP_CFP	1 bit	<p>Contention Free Period Capability</p> <p>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.</p> <p>It shall be set by SN in REG_REQ and by BN in REG_RSP.</p>
REG.CAP_DC	1 bit	<p>Direct Connection Capability</p> <p>1 if the device is able to perform direct connections; 0 if the device is not able to perform direct connections.</p> <p>It shall be set by SN in REG_REQ.</p>

Name	Length	Description
REG.ALV_F	1 bit	<p>Bit to indicate which ALV mechanism is required to be used by the new Service Node while it is part of this Subnetwork.</p> <p>1 ALV procedure of v1.4 shall be used;</p> <p>0 ALV procedure of v1.3.6 (section K.2.5) shall be used.</p> <p>It shall be set by BN in REG_RSP.</p> <p>Note: Base Node shall not selectively use different values of this bit between different Service Nodes in its Subnetwork. In case ALV procedure of v1.3.6 is used, all Service Nodes shall be instructed with REG.ALV_F bit set to 0.</p>
Reserved	1 bit	Always 0 for this version of the specification. Reserved for future use.
REG.CAP_ARQ	1 bit	<p>ARQ Capable</p> <p>1 if the device is able to establish ARQ connections;</p> <p>0 if the device is not able to establish ARQ connections.</p> <p>It shall be set by SN in REG_REQ and by BN in REG_RSP.</p>
REG.TIME	3 bits	<p>Time to wait for an ALV procedure before assuming the Service Node has been unregistered by the Base Node.</p> <p>ALV.TIME = 0 => 128 seconds ~ 2.1 minutes;</p> <p>ALV.TIME = 1 => 256 seconds ~ 4.2 minutes;</p> <p>ALV.TIME = 2 => 512 seconds ~ 8.5 minutes;</p> <p>ALV.TIME = 3 => 2048 seconds ~ 34.1 minutes;</p> <p>ALV.TIME = 4 => 4096 seconds ~ 68.3 minutes;</p> <p>ALV.TIME = 5 => 8192 seconds ~ 136.5 minutes;</p> <p>ALV.TIME = 6 => 16384 seconds ~ 273.1 minutes;</p> <p>ALV.TIME = 7 => 32768 seconds ~ 546.1 minutes;</p> <p>It shall be set by BN in REG_RSP.</p>

Name	Length	Description
REG.EUI-48	48 bit	EUI-48 of the Node EUI-48 of the Node requesting the Registration.
REG.RM_F	2 bits	Forces an encoding for the given node disabling robustness-management for its transmission, it can be disabled. Only used in REG_RSP. In all other message variants, this shall be 0. 0 - Disable, automatic robustness-management by the service nodes 1 - DBPSK_CC, device shall transmit always in DBPSK_CC 2 - DQPSK_R, device shall transmit always in DQPSK_R 3 - DBPSK_R, device shall transmit always in DBPSK_R
REG.SAR_SIZE	3 bits	Maximum SAR segment size the service node shall use. Only used in REG_RSP. In all other message variants, this shall be 0. 0: Not mandated by BN (SAR operates normally) 1: SAR = 16 bytes 2: SAR =32 bytes 3: SAR = 48 bytes 4: SAR =64 bytes 5: SAR =128 bytes 6: SAR =192 bytes 7: SAR =255 bytes
Reserved	3 bits	Always 0 for this version of the specification. Reserved for future use.
REG.CNT	32 bits	A counter to be used as the nonce for the registration PDU-s authentication/encryption.
REG.SWK	192 bits	Subnetwork key wrapped with KWK that shall be used to derive the Subnetwork working key

Name	Length	Description
REG.WK	192 bits	Encrypted authentication key wrapped with KWK. This is a random sequence meant to act as authentication mechanism.

1963 The PKT.SID field is used in this control packet as the Switch where the Service Node is registering. The
 1964 PKT.LNID field is used in this control packet as the Local Node Identifier being assigned to the Service Node
 1965 during the registration process negotiation.

1966 The REG.CAP_PA field is used to indicate the packet aggregation capability as discussed in Section 4.3.7. In
 1967 the uplink direction, this field is an indication from the registering Terminal Node about its own capabilities.
 1968 For the Downlink response, the Base Node evaluates whether or not all the devices in the cascaded chain
 1969 from itself to this Terminal Node have packet-aggregation capability. If they do, the Base Node shall set
 1970 REG.CAP_PA=1; otherwise REG.CAP_PA=0.

MSB												
REG.N	REG.R	REG.SPC	REG.CAP_R	REG.CAP_BC	REG.CAP_SW	REG.CAP_PA	REG.CAP_CFP	REG.CAP_DC	REG.ALV_F	Res.	REG.CAP_ARQ	REG.TIME
		REG.EUI48[0]					REG.EUI48[1]					
		REG.EUI48[2]					REG.EUI48[3]					
		REG.EUI48[4]					REG.EUI48[5]					
REG.PRM_F		REG.SAR_SIZE	<i>Reserved</i>					REG.CNT[0]				
		REG.CNT[1]					REG.CNT[2]					
		REG.CNT[3]					REG.SWK[0]					
							REG.SWK[1..2]					
							REG.SWK[3..4]					
							REG.SWK[5..6]					
							REG.SWK[7..8]					
							REG.SWK[9..10]					
							REG.SWK[11..12]					
							REG.SWK[13..14]					
							REG.SWK[15..16]					
							REG.SWK[17..18]					
							REG.SWK[19..20]					
							REG.SWK[21..22]					
		REG.SWK[23]					REG.WK[0]					
							REG.WK[1..2]					
							REG.WK[3..4]					
							REG.WK[5..6]					
							REG.WK[7..8]					
							REG.WK[9..10]					
							REG.WK[11..12]					
							REG.WK[13..14]					
							REG.WK[15..16]					
							REG.WK[17..18]					
							REG.WK[19..20]					
							REG.WK[21..22]					
							REG.WK[23]					

LSB

1971

1972

Figure 54 - REG control packet structure

Table 21 - 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 shall be lost; If R=1 any previous connection from this Node shall be maintained.
REG_RSP	1	< 0x3FFF	0	R	Registration response. This packet assigns the PKT.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.

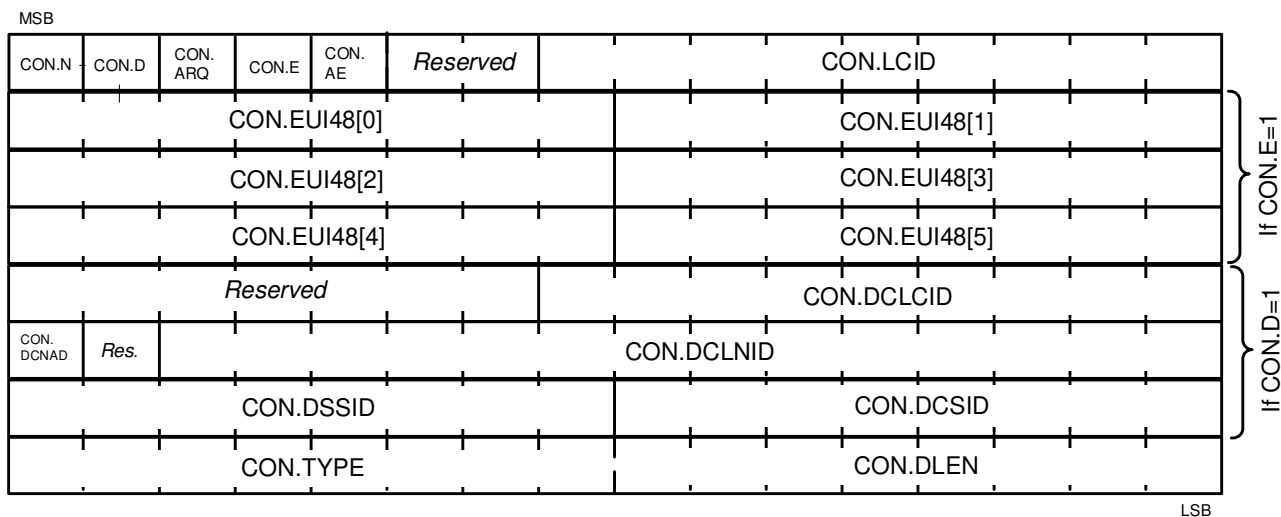
1974 Fields REG.SWK and REG.WK are of significance only for REG_RSP messages with Security Profiles 1 and 2
1975 (REG.SCP=1 and REG.SCP=2). For all other message-exchange variants using the REG control packet, these
1976 fields shall not be present reducing the length of payload.

1977 In REG_RSP message, the REG.SWK and REG.WK shall always be inserted wrapped with KWK.

1978 Field REG.CNT is of significance only for REG_REQ, REG_RSP and REG_REJ messages with Security Profiles 1
1979 and 2 (REG.SCP=1 and REG.SCP=2). For all other message-exchange variants using the REG control packet,
1980 these fields shall not be present reducing the length of payload.

1981 4.4.2.6.4 CON control packet (PKT.CTYPE = 2)

1982 This control packet is used for negotiating the connections. The description of the fields of this packet is given
1983 in Table 22 and Figure 55 The meaning of the packet differs depending on the direction of the packet and on
1984 the values of the different types. Table 23 shows the different interpretation of the packets. The packets are
1985 used during the connection establishment and closing.



1986
1987

Figure 55 - CON control packet structure

1988 Note that Figure 55 shows the complete message with all optional parts. When CON.D is 0, CON.DCNAD,
1989 CON.DSSID, CON.DCLNID, CON.DCLID, CON.DCSID and the reserved field between CON.DCNAD and
1990 CON.DSSID shall not be present in the message. Thus, the message shall be 6 octets smaller. Similarly, when
1991 CON.E is zero, the field CON.EUI-48 shall not be present, making the message 6 octets smaller.

1992

Table 22 - 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.
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.

Name	Length	Description
CON.AE	1 bit	<p>Use authentication and encryption on the data sent using this connection.</p> <ul style="list-style-type: none"> • CON.AE = 1 to encrypt and authenticate the data packets. • CON.AE = 0 to send the plain data without any authentication and encryption. <p>This flag is valid in both directions; each side shall set it to the desired value. The highest security will be applied in case the values do not match. So if any side sets this flag to 1 the data shall be authenticated and encrypted.</p> <p>NOTE: This flag only can be 1 on security profiles 1 and 2.</p>
<i>Reserved</i>	2 bits	<p>Reserved for future version of the protocol.</p> <p>This shall be 0 for this version of the protocol.</p>
CON.LCID	9 bits	<p>Local Connection Identifier.</p> <p>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.</p> <p>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.</p>
CON.EUI-48	48 bits (Present if CON.E=1)	<p>EUI-48 of destination/source Service Node/Base Node for connection request.</p> <p>When not performing a directed connection, this field shall not be included. When performing a directed connection, it may contain the SNA, indicating that the Base Node Convergence layer shall determine the EUI-48.</p> <ul style="list-style-type: none"> • CON.D = 0, Destination EUI-48; • CON.D = 1, Source EUI-48.
<i>Reserved</i>	7 bits (Present if CON.D=1)	<p>Reserved for future version of the protocol.</p> <p>This shall be 0 for this version of the protocol.</p>
CON.DCLCID	9 bits (Present if CON.D=1)	<p>Direct Connection LCID</p> <p>This field represents the LCID of the connection identifier to which the one being established shall be directly switched.</p>

Name	Length	Description
CON.DCNAD	1 bit (Present if CON.D=1)	Reserved for future version of the protocol. Direct Connection Not Aggregated at Destination This field represents the content of the PKT.NAD field after a direct connection Switch operation.
<i>Reserved</i>	1 bits (Present if CON.D=1)	Reserved for future version of the protocol. 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 shall 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.
CON.TYPE	8 bits	Connection type. The connection type (see Annex E) 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) (Present if CON.DLEN>0)	Connection specific parameters. These connections specific parameters are Convergence layer specific. They shall 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.

1993

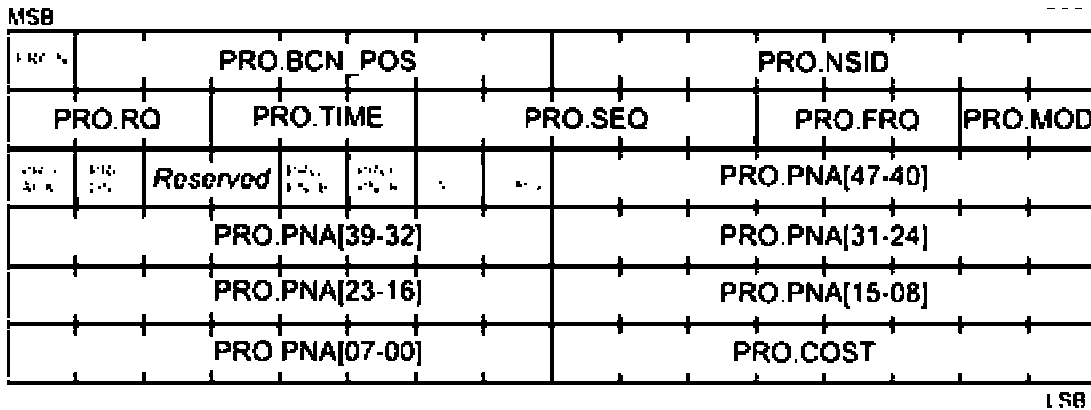
Table 23 - CON control packet types

Name	HDR.DO	CON.N	Description
CON_REQ_S	0	0	Connection establishment request initiated by the Service Node.


Name	HDR.DO	CON.N	Description
CON_REQ_B	1	0	The Base Node shall 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 shall 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.

1994 **4.4.2.6.5 PRO control packet (PKT.CTYPE = 3)**

1995 This control packet is used to promote a Service Node from Terminal function to Switch function. This control
1996 packet is also used to exchange information that is further used by the Switch Node to transmit its beacon.
1997 The description of the fields of this packet is given in Table 24, and Figure 57. The meaning of the packet
1998 differs depending on the direction of the packet and on the values of the different types.



1999
2000 **Figure 56 - PRO_REQ_S control packet structure**



2001
2002 **Figure 57 - PRO control packet structure**

2003 Note that Figure 56 includes all fields as used by a PRO_REQ_S message. All other messages are much smaller,
 2004 containing only PRO.N, PRO.RC, PRO.TIME and PRO.NSID as shown in Figure 57.

2005 **Table 24 - 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
PRO.BCN_POS	7 bits	Position of this beacon in symbols from the beginning of the frame.
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.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 that is being used by the terminal that shall become a switch. On a reception of this time in a PRO_REQ_B the Service Node shall reset the Keep-Alive timer in the same way as receiving an ALV_REQ_B.
PRO.SEQ	5 bits	The Beacon Sequence number when the specified change takes effect.
PRO.FRQ	3 bits	Transmission frequency of Beacon, 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 => <i>Reserved</i> FRQ = 7 => <i>Reserved</i>

PRO.MOD	2 bits	Modulation of the transmitted Beacons, encoded as: ENC = 0 => DBPSK + Convolutional Code ENC = 1 => Robust DQPSK ENC = 2 => Robust DBPSK ENC = 3 => Reserved
PRO.ACK	1 bit	Flag to differentiate the PRO_REQ_S from the PRO_ACK
PRO.DS	1 bit	Double switch flag. Used for switches that have to send a second beacon. This field is described in more detail in section 4.6.3.
<i>Reserved</i>	2 bits	Reserved for future versions of the protocol. Shall be set to 0 for this version of the protocol.
PRO.PN_BC	1 bit	Backwards Compatibility mode of the node represented by PRO.PNA. 1 if the device is backwards compatible with 1.3.6 PRIME 0 if it is not
PRO.PN_R	1 bit	Robust mode compatibility of the node represented by PRO.PNA. 1 if the device supports robust mode 0 if it is not
PRO.SWC_DC	1 bit	Direct Connection Switching Capability 1 if the device is able to behave as Direct Switch in direct connections. 0 otherwise
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.
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.COST	0 or 8 bits	Total 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 25 - PRO control packet types


Name	HDR. DO	PRO. N	PRO. ACK	PRO. NSID	Description
PRO_REQ_S	0	0	0	-	Used by terminal nodes to request a promotion to switch nodes. This is not part of any procedure, just an information message, so there shall not be any PRO_ACK/PRO_NACK in response. Used by switch nodes to request a beacon modulation change. In this case it is a procedure, so the base node shall respond with a PRO_ACK/PRO_NACK.
PRO_REQ_B	1	0	0	< 0xFF	The Base Node shall consider that the Service Node has promoted with the identifier PRO.NSID. <ul style="list-style-type: none"> To a terminal node: Promotion acceptance with allocating LSID or Promotion request initiated by the Base Node. To a switch node: Beacon information change initiated by the Base Node.
PRO_ACK	-	0	1	< 0xFF	Acknowledge. Used by both the Base Node and the Service Node to acknowledge with a positive answer the procedure. Procedures this message applies to are: PRO_REQ_S for beacon change modulation or PRO_REQ_B.
PRO_NACK	-	1	1	< 0xFF	Negative Acknowledge. Used by both the Base Node and the Service Node to acknowledge with a negative answer the procedure. Procedures this message applies to be: beacon change modulation.
PRO_DEM_S	0	1	0	< 0xFF	Used by Service Nodes to request a demotion, to reject a promotion or to positively acknowledge a demotion.
PRO_DEM_B	1	1	0	-	Used by the Base Node to request a demotion, to reject a promotion or to positively acknowledge a demotion.

2007 Table 25 shows the different interpretation of the packets. The promotion process is explained in more detail
2008 in 4.6.3.

2009 4.4.2.6.6 FRA control packet (PKT.CTYPE = 5)

2010 This control packet is broadcast from the Base Node and relayed by all Switch Nodes to the entire
2011 Subnetwork. It is used to circulate information on the change of Frame structure at a specific time in future.
2012 The description of fields of this packet is given in

2013 Table 26 and Figure 58.



2014
2015
2016

Figure 58 - FRA control packet structure


Table 26 - FRA control packet fields

Name	Length	Description
<i>Reserved</i>	2 bits	Reserved bits. Shall be set to 0b10 for this version of this specification
FRA.LEN	2 bits	Length of the frame to be applied in the next superframe. This shall be the <i>macFrameLength</i> , encoded with same semantics as the PIB attribute.
<i>FRA.PHYB</i> <i>C</i>	1 bit	The network is working on PHY backwards compatibility mode, all the nodes that need to send Type B PHY Frames shall use PHY backwards compatible frames. <ul style="list-style-type: none"> 0 if the subnet is not working in PHY backwards compatibility mode. 1 if the subnet is working in PHY backwards compatibility mode.
<i>Reserved</i>	1 bit	Reserved for future version of this protocol. In this version, this field shall be initialized to 0.
FRA.CFP_ LEN	10 bits	Length of CFP.
FRA.SEQ	5 bits	The Beacon Sequence number when the specified change takes effect.
<i>Reserved</i>	3 bits	Reserved for future version of this protocol. In this version this field shall be set to 0.

2017 **4.4.2.6.7 CFP control packet (PKT.CTYPE = 6)**

2018 This control packet is used for dedicated contention-free channel access time allocation to individual
2019 Terminal or Switch Nodes. The description of the fields of this packet is given in

2020 Table 27 and Figure 59. The meaning of the packet differs depending on the direction of the packet and on
2021 the values of the different types. Table 28 represents the different interpretation of the packets.



2022
2023

Figure 59 - CFP control packet structure

2024

Table 27 - CFP control message fields

Name	Length	Description
CFP.N	1 bit	0: denial of allocation/deallocation request; 1: acceptance of allocation/deallocation 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.
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.

2025

Table 28 - 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.

2026

4.4.2.6.8 ALV control packet (PKT.CTYPE = 7)

2027

2028

2029

2030

The ALV control message is used for Keep-Alive signaling between a Service Node, the Service Nodes above it and the Base Node. It is also used to test every hop in the path of that particular node performing robustness-management. Structures of these messages are shown in Figure 60, Figure 61 and Figure 62 and individual fields are enumerated in


2031

Table 29. The different Keep-Alive message types are shown in appropriate ALV.VALU(*) to zero.

2032

Table 30. These messages are sent periodically, as described in section 4.6.5.

2033




2034

2035

Figure 60 - ALV_RSP_S / ALV_REQ_B Control packet structure

2036

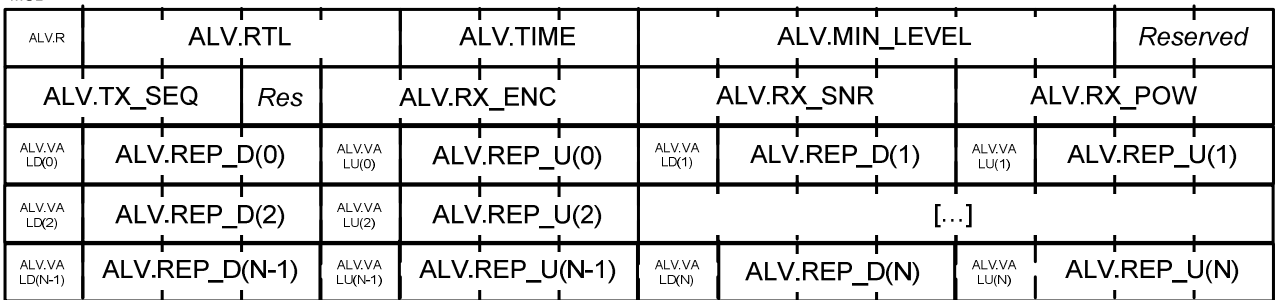


2037

2038

Figure 61 - ALV_ACK_B/ALV_ACK_S control packet structure

2039



2040

2041

Figure 62 - ALV_RSP_ACK Control packet structure

2042

Table 29 - ALV control message fields

Name	Length	Description
ALV.R	1 bit	Request/Response field. <ul style="list-style-type: none"> • 1 in the requests/response; • 0 in the acknowledges.
ALV.RTL	4 bits	Total number of repetitions left across the entire path.
ALV.TIME	3 bits	Time to wait for an ALV procedure before assuming the Service Node has been unregistered by the Base Node. <p>ALV.TIME = 0 => 128 seconds ~ 2.1 minutes;</p> <p>ALV.TIME = 1 => 256 seconds ~ 4.2 minutes;</p> <p>ALV.TIME = 2 => 512 seconds ~ 8.5 minutes;</p> <p>ALV.TIME = 3 => 2048 seconds ~ 34.1 minutes;</p> <p>ALV.TIME = 4 => 4096 seconds ~ 68.3 minutes;</p> <p>ALV.TIME = 5 => 8192 seconds ~ 136.5 minutes;</p> <p>ALV.TIME = 6 => 16384 seconds ~ 273.1 minutes;</p> <p>ALV.TIME = 7 => 32768 seconds ~ 546.1 minutes.</p>

ALV.MIN_LEVEL	6 bits	Minimum level of the node to add independent records to the REP_D/REP_U table. For downlink case, if the switch's level+1 is equal or lower than ALV.MIN_LEVEL it shall sum its repetitions to the first record (record number 0: ALV.REP_U(0)). For uplink case, if the switch's (or terminal's) level is equal or lower than ALV.MIN_LEVEL it shall sum its repetitions to the first record.
Reserved	2 bit	Reserved for future use. Shall be 0 for this version of the specification.
ALV.TX_SEQ	3 bits	Sequence of number of transmissions, to keep track if the loss in the ALV process is due to the REQ/RSP process or the ACK. This is to avoid an incorrect evaluation of downlink/uplink because of ACK loses. All the ALV operations shall start with sequence number 0, every time a node starts a hop level operation it shall set this field to 0, and each repetition it shall increase it until ACK is received.
ALV.VALD(*)	1 bit	Flag to indicate that the REP_D record contains valid information. <ul style="list-style-type: none"> • 1 information contained in REP_D records is valid; • 0 information in REP_D shall be discarded.
ALV.REP_D(*)	3 bits	Number of repetitions for the given downlink hop. Valid values: <ul style="list-style-type: none"> • 0-5 : Number of repetitions; • 6 : 6 or more repetitions (for record as a sum of various levels); • 7 : All the retries finished for this hop.
ALV.VALU(*)	1 bit	Flag to indicate that the REP_U record contains valid information. <ul style="list-style-type: none"> • 1 information contained in REP_D records is valid; • 0 information in REP_D shall be discarded.
ALV.REP_U(*)	3 bits	Number of repetitions for the given uplink hop. In the ALV_REQ_B the base node shall fill these fields with the repetitions of each hop for the last ALV procedure of that hop. In the ALV_RSP_S the service node shall fill the field with the uplink repetitions. Valid values: <ul style="list-style-type: none"> • 0-5 : Number of repetitions; • 6 : 6 or more repetitions (for record as a sum of various levels); • 7 : All the retries finished for this hop.
ALV.RX_SNR	4 bits	Signal to Noise Ratio at which the ALV_REQ_B/ALV_RSP_S was received.
ALV.RX_POW	4 bits	Power at which the ALV_REQ_B/ALV_RSP_S was received.
ALV.RX_ENC	4 bits	Encoding at which the ALV_REQ_B/ALV_RSP_S was received. <ul style="list-style-type: none"> • 0 – DBPSK • 1 – DQPSK • 2 – D8PSK • 3 – Not used • 4 – DBPSK + Convolutional Code

		<ul style="list-style-type: none"> • 5 – DQPSK + Convolutional Code • 6 – D8PSK + Convolutional Code • 7-11 – Not used • 12 – Robust DBPSK • 13 – Robust DQPSK • 14 – Not used • 15 – Outdated information
--	--	---

2040 The * symbol means that there are a variable number of records for the same field, each record shall be
2041 fulfilled by a Service Node in the path to the Terminal Node that shall receive the ALV, Position N shall be
2042 the hop of the Service Node target of the ALV procedure, N-1 shall be the parent Switch Node of that node,
2043 and so on. For the switches with level below or equal than ALV.MIN_LEVEL shall add their repetitions
2044 information to the record ALV.REP_U(0)/ALV.REP_D(0). The Base Node shall make sure that the number of
2045 records is correct for the given ALV.MIN_LEVEL value.

2046 The base node shall fill the ALV.REP_U(*) registries with the last ALV operation’s uplink retries for each hop,
2047 if it does not have that information it shall reset the appropriate ALV.VALU(*) to zero.

2048 **Table 30 - Keep-Alive control packet types**


Name	HDR.DO	ALV.R	Description
ALV_REQ_B	1	1	Keep-Alive request message.
ALV_ACK_B	1	0	Keep-Alive acknowledge to a response.
ALV_RSP_S	0	1	Keep-Alive response message in case the node is not the target node (PKT.SID != receiver SSID)
ALV_RSP_ACK	0	1	Keep-Alive response acknowledge message in case the node is the target node (PKT.SID == receiver SSID)
ALV_ACK_S	0	0	Keep-Alive acknowledge to a request.

2049 **4.4.2.6.9 MUL control packet (PKT.CTYPE = 8)**

2050 The MUL message is used to control multicast group membership. The structure of this message and the
2051 meanings of the fields are described in Table 31 and

2052 Figure 63. The message can be used in different ways as described in

2053 Table 32.



2054

2055

Figure 63 - MUL control packet structure

Table 31 - MUL control message fields

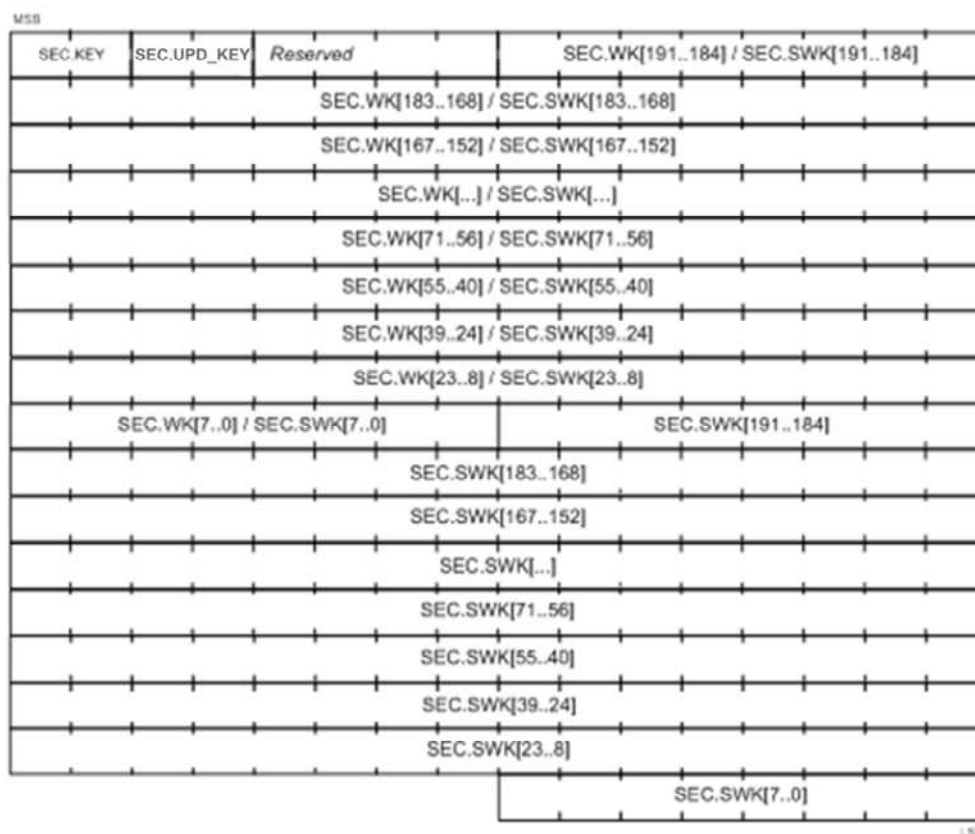
Name	Length	Description
MUL.N	1 bit	<p>Negative</p> <ul style="list-style-type: none"> MUL.N = 1 for the negative multicast connection, i.e. multicast group leave. MUL.N = 0 for the positive multicast connection, i.e. multicast group join.
MUL.AE	1 bit	<p>Use authentication and encryption on the data sent using this multicast connection.</p> <ul style="list-style-type: none"> MUL.AE = 1 to encrypt and authenticate the data packets. MUL.AE = 0 to send the plain data without any authentication and encryption. <p>This flag is valid in both directions; each side shall set it to the desired value. The highest security will be applied in case the values do not match. So if any side sets this flag to 1 the data shall be authenticated and encrypted.</p> <p>NOTE: This flag only can be 1 on security profiles 1 and 2.</p>
Reserved	5 bits	Reserved for future version of the protocol. 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 E.
MUL.DLEN	8 bits	Length of data in bytes in the MUL.DATA field
MUL.DATA	(variable)(Present if MUL.DLEN >0)	Connection specific parameters. These connections specific parameters are Convergence layer specific. They shall 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 32 - MUL control message types

Name	HDR.DO	MUL.N	PKT.LNID	Description
MUL_JOIN_S	0	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	>0	The Base Node shall 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	>0	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	>0	The Base Node shall 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.
MUL_SW_LEAVE_B	1	1	0	The switch node shall stop switching multicast data for the multicast group. This message is always initiated by the base node. The addressing shall be with the switch's SSID and LNID == 0 to distinguish this message from MUL_LEAVE_B.
MUL_SW_LEAVE_S	0	1	0	The switch node is no longer switching multicast data for the multicast group. This message is sent as a response to MUL_SW_LEAVE_B. The addressing shall be with the switch's SSID and LNID == 0 to distinguish this message from MUL_LEAVE_S.

2059 **4.4.2.6.10 SEC control packet (PKT.CTYPE = 10)**

2060 The SEC control message is a unicast message transmitted authenticated and encrypted (WK) by the Base
 2061 Node to every node in the Subnetwork to update the WK and SWK. The random sequence used by devices
 2062 in a Subnetwork is dynamic and changes from time to time to ensure a robust security framework. The
 2063 structure of this message is shown in Table 33 and Figure 64. Further details of security mechanisms are given
 2064 in Section 4.3.8.



2065
2066 **Figure 64 - SEC control packet structure**

2067 **Table 33 - SEC control message fields**

Name	Length	Description
SEC.KEY	2 bits	In the SEC_REQ it indicates which key is present In the SEC_RSP it indicates which key was to be updated 0 - reserved 1 – only SEC.WK is present / only SEC.WK to be updated 2 – only SEC.SWK is present / only SEC.SWK to be updated 3 – SEC.WK and SEC.SWK are present / SEC.WK and SEC.SWK to be updated

Name	Length	Description
SEC.UPDA TED_KEY	2 bits	In the SEC_RSP it indicates which key has being updated 0 - reserved 1 – only SEC.WK was updated 2 – only SEC.SWK was updated 3 – SEC.WK and SEC.SWK were updated Only used in SEC_RSP. In all other message variants, this shall be 0.
Reserved	4 bits	Shall always be encoded as 0 in this version of the specification.
SEC.WK	192 bits	(optional in SEC_REQ, not present in SEC_RSP) Working Key wrapped by KWK.
SEC.SWK	192 bits	(optional in SEC_REQ, not present in SEC_RSP) Subnetwork Working Key wrapped by KWK.

2068

Table 34 - SEC control packet types

Name	HDR.DO	Description
SEC_REQ	1	Security key update request message.
SEC_RSP	0	Security key update acknowledge to a request.

2069

4.4.3 Promotion Needed PDU

2070

If a Node is Disconnected and it does not have connectivity with any existing Switch Node, it shall send notifications to its neighbors to indicate the need for the promotion of any available Terminal Node.

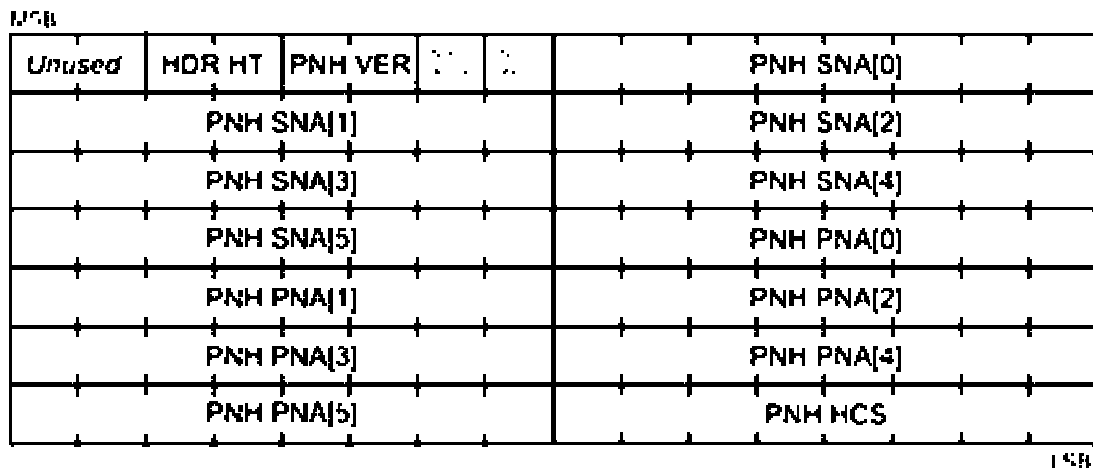
2071

2072

Figure 39 represents the Promotion Needed MAC PDU (PNPDU) that must be sent on an irregular basis in this situation.

2073

2074



2075

2076

Figure 65 - Promotion Need MAC PDU

2077 Table 35 shows the promotion need MAC PDU fields.

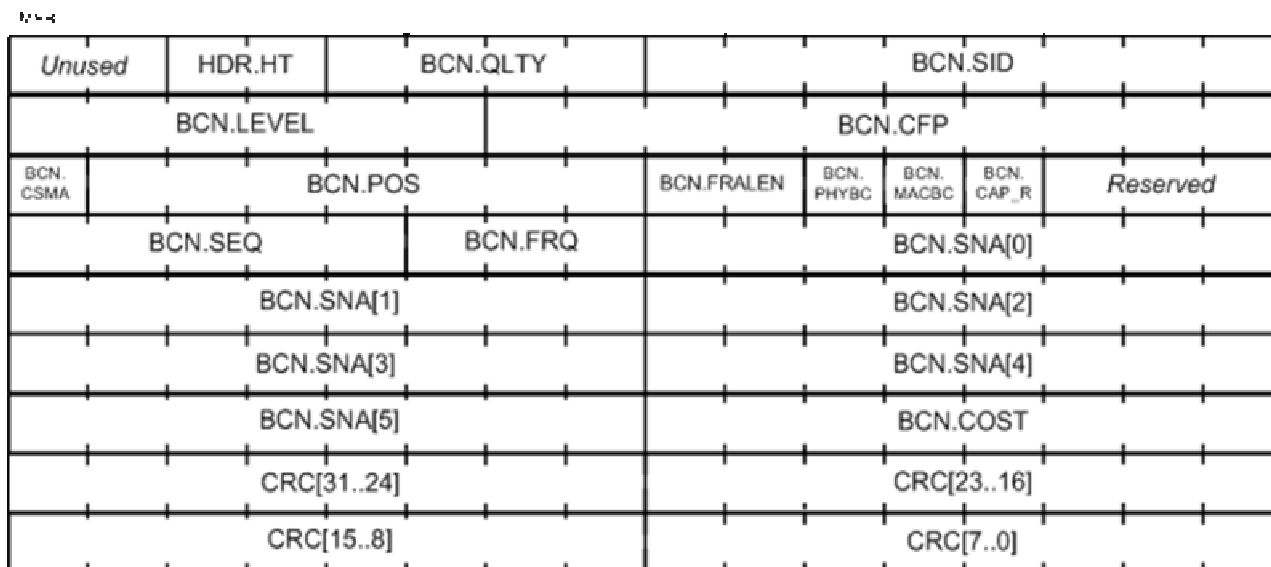
2078 **Table 35 - Promotion Need MAC PDU fields**

Name	Length	Description
<i>Unused</i>	2 bits	Unused bits which are always 0; included for alignment with MAC_H field in PPDU header (Section 3.3.3).
HDR.HT	2 bits	Header Type HDR.HT = 1 for the Promotion Need MAC PDU
PNH.VER	2 bits	Version of PRIME Specification: <ul style="list-style-type: none"> • 0 – 1.3.6 PRIME • 1 – 1.4 PRIME • 2,3 – Reserved for future use
PNH.CAP_ R	1 bit	Flag to define if the node supports Robust mode <ul style="list-style-type: none"> • 0 – If the node does not support Robust mode • 1 – if the node does support Robust mode
PNH.CAP_ BC	1 bit	Flag to define if the node supports Backwards Compatibility with 1.3.6 version of the specification <ul style="list-style-type: none"> • 0 – The node does not support Backwards Compatibility with 1.3.6 • 1 – The node supports Backwards Compatibility with 1.3.6
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 $M(x)$.

2079 As it is always transmitted by unsynchronized Nodes and, therefore, prone to creating collisions, it is a special
 2080 reduced size header.

2081 4.4.4 Beacon PDU

2082 Beacon PDU (BPDU) is transmitted by every Switch device on the Subnetwork, including the Base Node. The
 2083 purpose of this PDU is to circulate information on MAC frame structure and therefore channel access to all
 2084 devices that are part of this Subnetwork. The BPDU is transmitted at definite fixed intervals of time and is
 2085 also used as a synchronization mechanism by Service Nodes. Figure 66 below shows contents of a beacon
 2086 transmitted by the Base Node and each Switch Device.



2087

2088

Figure 66 - Beacon PDU structure

2089 Table 36 shows the beacon PDU fields.

2090

Table 36 - Beacon PDU fields

Name	Length	Description
Unused	2 bits	Unused bits which are always 0; included for alignment with MAC_H field in PPDU header (Fig 7, Section 3.3.3).
HDR.HT	2 bits	Header Type HDR.HT = 2 for Beacon PDU

Name	Length	Description
BCN.QLTY	4 bits	<p>Quality of round-trip connectivity from this Switch Node to the Base Node. ,with the following meaning:</p> <p>BCN.QLTY = 0 if $1/2 < \text{rtdp} \leq 1$</p> <p>BCN.QLTY = 1 if $3/8 < \text{rtdp} \leq 1/2$</p> <p>BCN.QLTY = 2 if $1/4 < \text{rtdp} \leq 3/8$</p> <p>BCN.QLTY = 3 if $3/16 < \text{rtdp} \leq 1/4$</p> <p>BCN.QLTY = 4 if $1/8 < \text{rtdp} \leq 3/16$</p> <p>BCN.QLTY = 5 if $3/32 < \text{rtdp} \leq 1/8$</p> <p>BCN.QLTY = 6 if $1/16 < \text{rtdp} \leq 3/32$</p> <p>BCN.QLTY = 7 if $3/64 < \text{rtdp} \leq 1/16$</p> <p>BCN.QLTY = 8 if $1/32 < \text{rtdp} \leq 3/64$</p> <p>BCN.QLTY = 9 if $3/128 < \text{rtdp} \leq 1/32$</p> <p>BCN.QLTY = 10 if $1/64 < \text{rtdp} \leq 3/128$</p> <p>BCN.QLTY = 11 if $3/256 < \text{rtdp} \leq 1/64$</p> <p>BCN.QLTY = 12 if $1/128 < \text{rtdp} \leq 3/256$</p> <p>BCN.QLTY = 13 if $3/512 < \text{rtdp} \leq 1/128$</p> <p>BCN.QLTY = 14 if $1/256 < \text{rtdp} \leq 3/512$</p> <p>BCN.QLTY = 15 if $\text{rtdp} \leq 1/256$</p> <p>where:</p> <p>rtdp = Round Trip Drop Probability. Probability for a packet to be dropped when it is supposed to go downlink and be answered uplink (or the other way around) between the Base Node and the Switch Node.</p> <p>It is up to the manufacturer how to detect it, and It doesn't have to be very accurate, just an estimation. As a guideline, ALV packets can be used to calculate this field.</p>
BCN.SID	8 bits	Switch identifier of transmitting Switch
BCN.LEVEL	6 bits	Hierarchy of transmitting Switch in Subnetwork
BCN.CFP	10 bits	CFP length in symbols.
BCN.CSMA	1 bit	<p>CSMA/CA Algorithm used in the Subnetwork.</p> <ul style="list-style-type: none"> • 0 – CSMA/CA Algorithm 1 (i.e. v1.4) • 1 – CSMA/CA Algorithm 2 (i.e. v1.3.6, as in backward compatibility mode)

Name	Length	Description
BCN.POS	7 bits	Position of this beacon in symbols from the beginning of the frame.
BCN.FRA_LEN	2 bits	Length of the frame. <ul style="list-style-type: none"> • 0 - 276 symbols • 1 - 552 symbols • 2 - 828 symbols • 3 - 1104 symbols
BCN.PHYBC	1 bit	PHY backwards compatibility mode: 0 – The network is working in normal mode. 1 - The network is working on PHY backwards compatibility mode, all the nodes that need to send Type B PHY Frames shall use PHY backwards compatible frames.
BCN.MACBC	1 bit	MAC backward compatibility mode: 0 – The network is working in 1.4 mode. 1 – The network is working in MAC backward compatibility mode, see section 4.9 for details.
BCN.CAP_R	1 bit	Robust Mode Capable 1 – The device is able to transmit/receive robust mode frames. 0 – The device is not able to transmit/receive robust mode frames.
Reserved	3 bits	Always 0 for this version of the specification. Reserved for future use.
BCN.SEQ	5 bits	Sequence number of this BPDU in super frame. Incremented for every beacon the Base Node sends and is propagated by Switch through its BPDU such that 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 7 = Reserved
BCN.SNA	48 bits	Subnetwork identifier in which the Switch transmitting this BPDU is located

Name	Length	Description
BCN.COST	8 bits	<p>Total cost from the transmitting Switch Node to the Base Node. The cost of a single hop is calculated based on modulation scheme used on that hop in both downlink and uplink direction. Values are derived as follows:</p> <p style="margin-left: 40px;">8PSK = 0 QPSK = 0 BPSK = 0 8PSK_F = 0 QPSK_F = 1 BPSK_F = 2 QPSK_R = 4 BPSK_R = 8</p> <p>The Base Node shall transmit in its beacon a BCN.COST of 0. A Switch Node shall transmit in its beacon the value of BCN.COST received from its upstream Switch Node, plus the cost of the upstream hop to its upstream Switch, calculated as the addition of both uplink and downlink costs. When this value is larger than what can be held in BCN.UPCOST the maximum value of BCN.COST shall be used.</p>
CRC	32 bits	<p>The CRC shall be calculated with the same algorithm as the one defined for the CRC field of the MAC PDU (see section 0 for details). For CRC calculation the field CRC is set to the constant 0x00010400. The CRC shall be calculated over the whole BPDU, including constant CRC field</p>

2091 The BPDU is also used to detect when the uplink Switch is no longer available either by a change in the
2092 characteristics of the medium or because of failure etc. If a Service Node fails to receive all the expected
2093 beacons during Nmiss-beacon superframes it shall declare the link to its Switch as unusable. The Service Node
2094 shall stop sending beacons itself if it is acting as a Switch. It shall close all existing MAC connections. The
2095 Service Node then enters the initial Disconnected state and searches for a Subnetwork join. This mechanism
2096 complements the Keep-Alive mechanism which is used by a Base Node and its switches to determine when
2097 a Service Node is lost.

2098 **4.5 MAC Service Access Point**

2099 **4.5.1 General**

2100 The MAC service access point provides several primitives to allow the Convergence layer to interact with the
2101 MAC layer. This section aims to explain how the MAC may be used. An implementation of the MAC may not
2102 use all the primitives listed here; it may use other primitives; or it may have a function-call based interface
2103 rather than message-passing, etc. These are all implementation issues which are beyond the scope of this
2104 specification.

2105 The .request primitives are passed from the CL to the MAC to request the initiation of a service. The
 2106 .indication and .confirm primitives are passed from the MAC to the CL to indicate an internal MAC event that
 2107 is significant to the CL. This event may be logically related to a remote service request or may be caused by
 2108 an event internal to the local MAC. The .response primitive is passed from the CL to the MAC to provide a
 2109 response to a .indication primitive. Thus, the four primitives are used in pairs, the pair .request and .confirm
 2110 and the pair .indication and .response. This is shown in Figure 67, Figure 68, Figure 69 and Figure 70.




Figure 67 - Establishment of a Connection




Figure 68 - Failed establishment of a Connection




Figure 69 - Release of a Connection




Figure 70 - Transfer of Data

2111 Table 33 represents the list of available primitives in the MAC-SAP:

2112 **Table 37 - 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

Service Node primitives
MAC_RELEASE.confirm
MAC_JOIN.request
MAC_JOIN.Response
MAC_JOIN.indication
MAC_JOIN.confirm
MAC_LEAVE.request
MAC_LEAVE.indication
MAC_LEAVE.confirm
MAC_DATA.request
MAC_DATA.confirm
MAC_DATA.indication

Base Node primitives
MAC_RELEASE.confirm
MAC_JOIN.request
MAC_JOIN.response
MAC_JOIN.indication
MAC_JOIN.confirm
MAC_LEAVE.request
MAC_LEAVE.indication
MAC_LEAVE.confirm
MAC_REDIRECT.response
MAC_DATA.request
MAC_DATA.confirm
MAC_DATA.indication

2113 **4.5.2 Service Node and Base Node signalling primitives**

2114 **4.5.2.1 General**

2115 The following subsections describe primitives which are available in both the Service Node and Base Node
 2116 MAC-SAP. These are signaling primitives only and used for establishing and releasing MAC connections.

2117 **4.5.2.2 MAC_ESTABLISH**

2118 **4.5.2.2.1 General**

2119 The MAC_ESTABLISH primitives are used to manage a connection establishment.

2120 **4.5.2.2.2 MAC_ESTABLISH.request**

2121 The MAC_ESTABLISH.request primitive is passed to the MAC layer entity to request the connection
 2122 establishment.

2123 The semantics of this primitive are as follows:

2124 *MAC_ESTABLISH.request{EUI-48, Type, Data, DataLength, ARQ, CfBytes, AE}*

2125 The EUI-48 parameter of this primitive is used to specify the address of the Node to which this connection
 2126 will be addressed. The MAC will internally transfer this to an address used by the MAC layer. When the CL of
 2127 a Service Node wishes to connect to the Base Node, it uses the EUI-48 00:00:00:00:00:00. However, when
 2128 the CL of a Service Node wishes to connect to another Service Node on the Subnetwork, it uses the EUI-48 of

2129 that Service Node. This will then trigger a direct connection establishment. However, whether a normal or a
2130 directed connection is established is transparent to the Service Node MAC SAP. As the EUI-48 of the Base
2131 Node is the SNA, the connection could also be requested from the Base Node using the SNA.

2132 The *Type* parameter is an identifier used to define the type of the Convergence layer that should be used for
2133 this connection (see Annex E). This parameter is 1 byte long and will be transmitted in the CON.TYPE field of
2134 the connection request.

2135 The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the local CL
2136 and the remote CL. This parameter will be transmitted in the CON.DATA field of the connection request.

2137 The *DataLength* parameter is the length of the *Data* parameter in bytes.

2138 The *ARQ* parameter indicates whether or not the ARQ mechanism should be used for this connection. It is a
2139 Boolean type with a value of true indicating that ARQ will be used.

2140 The *CfBytes* parameter is used to indicate whether or not the connection should use the contention or
2141 contention-free channel access scheme. When *CfBytes* is zero, contention-based access should be used.
2142 When *CfBytes* is not zero, it indicates how many bytes per frame should be allocated to the connection using
2143 CFP packets.

2144 The *AE* parameter indicates whether or not the information transmitted in this connection is encrypted. It is
2145 a Boolean type with a value of true indicating that encryption will be used.

2146 **4.5.2.2.3 MAC_ESTABLISH.indication**

2147 The MAC_ESTABLISH.indication is passed from the MAC layer to indicate that a connection establishment
2148 was initiated by a remote Node.

2149 The semantics of this primitive are as follows:

2150 *MAC_ESTABLISH.indication*{*ConHandle*, *EUI-48*, *Type*, *Data*, *DataLength*, *CfBytes*, *AE*}

2151 The *ConHandle* is a unique identifier interchanged to uniquely identify the connection being indicated. It has
2152 a valid meaning only in the MAC SAP, used to have a reference to this connection between different
2153 primitives.

2154 The *EUI-48* parameter indicates which device on the Subnetwork wishes to establish a connection.

2155 The *Type* parameter is an identifier used to define the type of the Convergence layer that should be used for
2156 this connection. This parameter is 1 byte long and it is received in the CON.TYPE field of the connection
2157 request.

2158 The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote
2159 CL and the local CL. This parameter is received in the CON.DATA field of the connection request.

2160 The *DataLength* parameter is the length of the *Data* parameter in bytes.

2161 The *CfBytes* parameter is used to indicate if the connection should use the contention or contention-free
2162 channel access scheme. When *CfBytes* is zero, contention-based access will be used. When *CfBytes* is not
2163 zero, it indicates how many bytes per frame the connection would like to be allocated.

2164 The *AE* parameter indicates whether or not the information transmitted in this connection is encrypted. It is
2165 a Boolean type with a value of true indicating that encryption will be used.

2166 4.5.2.2.4 MAC_ESTABLISH.response

2167 The MAC_ESTABLISH.response is passed to the MAC layer to respond with a MAC_ESTABLISH.indication.

2168 The semantics of this primitive are as follows:

2169 $MAC_ESTABLISH.response\{ConHandle, Answer, Data, DataLength, AE\}$

2170 The *ConHandle* parameter is the same as the one that was received in the MAC_ESTABLISH.indication.

2171 The *Answer* parameter is used to notify the MAC of the action to be taken for this connection establishment.
2172 This parameter may have one of the values in Table 38.

2173 The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote
2174 CL and the local CL. This parameter is received in the CON.DATA field of the connection response.

2175 The *DataLength* parameter is the length of the Data parameter in bytes.

2176 Data may be passed to the caller even when the connection is rejected, i.e. Answer has the value 1. The data
2177 may then optionally contain more information as to why the connection was rejected.

2178 The *AE* parameter indicates whether or not the information transmitted in this connection is encrypted. It is
2179 a Boolean type with a value of true indicating that encryption will be used.

2180 **Table 38 - Values of the *Answer* parameter in MAC_ESTABLISH.response primitive**

<i>Answer</i>	Description
<i>Accept</i> = 0	The connection establishment is accepted.
<i>Reject</i> = 1	The connection establishment is rejected.

2181 4.5.2.2.5 MAC_ESTABLISH.confirm

2182 The MAC_ESTABLISH.confirm is passed from the MAC layer as the remote answer to a
2183 MAC_ESTABLISH.request.

2184 The semantics of this primitive are as follows:

2185 $MAC_ESTABLISH.confirm\{ConHandle, Result, EUI-48, Type, Data, DataLength, AE\}$

2186 The *ConHandle* is a unique identifier to uniquely identify the connection being indicated. It has a valid
2187 meaning only in the MAC SAP, used to have a reference to this connection between different primitives. The
2188 value is only valid if the *Result* parameter is 0.

2189 The *Result* parameter indicates the result of the connection establishment process. It may have one of the
2190 values in Table 35 .

2191 The *EUI-48* parameter indicates which device on the Subnetwork accepted or refused to establish a
2192 connection.

2193 The *Type* parameter is an identifier used to define the type of the Convergence layer that should be used for
2194 this connection. This parameter is 1 byte long and it is received in the CON.TYPE field of the connection
2195 request

2196 The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote
2197 CL and the local CL. This parameter is received in the CON.DATA field of the connection response.

2198 The *DataLength* parameter is the length of the Data parameter in bytes.

2199 Data may be passed to the caller even when the connection is rejected, i.e. Result has the value 1. The data
2200 may then optionally contain more information as to why the connection was rejected.

2201 The *AE* parameter indicates whether or not the information transmitted in this connection is encrypted. It is
2202 a Boolean type with a value of true indicating that encryption will be used.

2203

Table 39 - Values of the *Result* parameter in MAC_ESTABLISH.confirm primitive

Result	Description
<i>Success = 0</i>	The connection establishment was successful.
<i>Reject = 1</i>	The connection establishment failed because it was rejected by the remote Node.
<i>Timeout = 2</i>	The connection establishment process timed out.
<i>No bandwidth = 3</i>	There is insufficient available bandwidth to accept this contention-free connection.
<i>No Such Device = 4</i>	A device with the destination address cannot be found.
<i>Redirect failed =5</i>	The Base Node attempted to perform a redirect which failed.
<i>Not Registered = 6</i>	The Service Node is not registered.
<i>No More LCIDs = 7</i>	All available LCIDs have been allocated.
<i>Unsupported SP = 14</i>	Device doesn't support SP > 0 but asked for an encrypted connection establishment.

2204 **4.5.2.3 MAC_RELEASE**

2205 **4.5.2.3.1 General**

2206 The MAC_RELEASE primitives are used to release a connection.

2207 **4.5.2.3.2 MAC_RELEASE.request**

2208 The MAC_RELEASE.request is a primitive used to initiate the release process of a connection.

2209 The semantics of this primitive are as follows:

2210 *MAC_RELEASE.request{ConHandle}*

2211 The *ConHandle* parameter specifies the connection to be released. This handle is the one that was obtained
2212 during the MAC_ESTABLISH primitives.

2213 **4.5.2.3.3 MAC_RELEASE.indication**

2214 The MAC_RELEASE.indication is a primitive used to indicate that a connection is being released. It may be
2215 released because of a remote operation or because of a connectivity problem.

2216 The semantics of this primitive are as follows:

2217 *MAC_RELEASE.indication{ConHandle, Reason}*

2218 The *ConHandle* parameter specifies the connection being released. This handle is the one that was obtained
2219 during the MAC_ESTABLISH primitives.

2220 The *Reason* parameter may have one of the values given in Table 36.

2221 **Table 40 - Values of the *Reason* parameter in MAC_RELEASE.indication primitive**

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

2222 **4.5.2.3.4 MAC_RELEASE.response**

2223 The MAC_RELEASE.response is a primitive used to respond to a connection release process.

2224 The semantics of this primitive are as follows:

2225 *MAC_RELEASE.response{ConHandle, Answer}*

2226 The *ConHandle* parameter specifies the connection being released. This handle is the one that was obtained
2227 during the MAC_ESTABLISH primitives.

2228 The *Answer* parameter may have one of the values given in Table 37 This parameter may not have the value
2229 "*Reject = 1*" because a connection release process cannot be rejected.

2230 **Table 41 - Values of the *Answer* parameter in MAC_RELEASE.response primitive**

<i>Answer</i>	Description
<i>Accept = 0</i>	The connection release is accepted.

2231 After sending the MAC_RELEASE.response the ConHandle is no longer valid and should not be used.

2232 **4.5.2.3.5 MAC_RELEASE.confirm**

2233 The MAC_RELEASE.confirm primitive is used to confirm that the connection release process has finished.

2234 The semantics of this primitive are as follows:

2235
$$\text{MAC_RELEASE.confirm}\{\text{ConHandle}, \text{Result}\}$$

2236 The *ConHandle* parameter specifies the connection released. This handle is the one that was obtained during
2237 the MAC_ESTABLISH primitives.

2238 The *Result* parameter may have one of the values given in Table 38

2239 **Table 42 - Values of the Result parameter in MAC_RELEASE.confirm primitive**

Result	Description
<i>Success = 0</i>	The connection release was successful.
<i>Timeout = 2</i>	The connection release process timed out.
<i>Not Registered = 6</i>	The Service Node is no longer registered.

2240 After the reception of the MAC_RELEASE.confirm the ConHandle is no longer valid and should not be used.

2241 **4.5.2.4 MAC_JOIN**

2242 **4.5.2.4.1 General**

2243 The MAC_JOIN primitives are used to join to a broadcast or multicast connection and allow the reception of
2244 such packets.

2245 **4.5.2.4.2 MAC_JOIN.request**

2246 The MAC_JOIN.request primitive is used:

- 2247
- 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 Base Node : to invite a Service Node to join a particular multicast group

2249

2250 Depending on which device makes the join-request, this SAP can have two different variants. First variant
2251 shall be used on Base Nodes and second on Service Nodes:

2252 The semantics of this primitive are as follows:

2253
$$\text{MAC_JOIN.request}\{\text{Broadcast}, \text{ConHandle}, \text{EUI-48}, \text{Type}, \text{Data}, \text{DataLength}, \text{AE}\}$$

2254
$$\text{MAC_JOIN.request}\{\text{Broadcast}, \text{Type}, \text{Data}, \text{DataLength}, \text{AE}\}$$

2255 The *Broadcast* parameter specifies whether the JOIN operation is being performed for a broadcast
2256 connection or for a multicast operation. It should be 1 for a broadcast operation and 0 for a multicast
2257 operation. In case of broadcast operation, EUI-48, Data, DataLength are not used.

2258 ConHandle indicates the handle to be used with for this multicast join. In case of first join request for a new
2259 multicast group, ConHandle will be set to 0. For any subsequent EUI additions to an existing multicast group,
2260 ConHandle will serve as index to respective multicast group.

2261 The EUI-48 parameter is used by the Base Node to specify the address of the Node to which this join request
2262 will be addressed. The MAC will internally transfer this to an address used by the MAC layer. When the CL of
2263 a Service Node initiates the request, it uses the EUI-48 00:00:00:00:00:00.

2264 The Type parameter defines the type of the Convergence layer that will send/receive the data packets. This
2265 parameter is 1 byte long and will be transmitted in the MUL.TYPE field of the join request.

2266 The Data parameter is a general purpose buffer to be interchanged for the negotiation between the remote
2267 CL and the local CL. This parameter is received in the MUL.DATA field of the connection request. In case the
2268 CL supports several multicast groups, this Data parameter will be used to uniquely identify the group

2269 The DataLength parameter is the length of the Data parameter in bytes.

2270 If Broadcast is 1, the MAC will immediately issue a MAC_JOIN.confirm primitive since it does not need to
2271 perform any end-to-end operation. For a multicast operation the MAC_JOIN.confirm is only sent once
2272 signaling with the uplink Service Node/Base Node is complete.

2273 The AE parameter indicates whether or not the information transmitted in this connection is encrypted. It is
2274 a Boolean type with a value of true indicating that encryption will be used.

2275 4.5.2.4.3 MAC_JOIN.confirm

2276 The MAC_JOIN.confirm primitive is received to confirm that the MAC_JOIN.request operation has finished.

2277 The semantics of this primitive are as follows:

2278 $MAC_JOIN.confirm\{ConHandle, Result, AE\}$

2279 The *ConHandle* is a unique identifier to uniquely identify the connection being indicated. It has a valid
2280 meaning only in the MAC SAP, used to have a reference to this connection between different primitives. The
2281 value is only valid if the *Result* parameter is 0. When the MAC receives packets on this connection, they will
2282 be passed upwards using the MAC_DATA.indication primitive with this *ConHandle*.

2283 The Result parameter indicates the result of multicast group join process. It may have one of the values given
2284 in Table 39.

2285 The AE parameter indicates whether or not the information transmitted in this connection is encrypted. It is
2286 a Boolean type with a value of true indicating that encryption will be used.

2287 **Table 43 - Values of the Result parameter in MAC_JOIN.confirm primitive**

Result	Description
Success = 0	The connection establishment was successful.

Result	Description
<i>Reject = 1</i>	The connection establishment failed because it was rejected by the upstream Service Node/Base Node.
<i>Timeout = 2</i>	The connection establishment process timed out.
<i>Unsupported SP = 14</i>	Device doesn't support SP >0 but asked to join a multicast group with an encrypted connection

2288 **4.5.2.4.4 MAC_JOIN.indication**

2289 On the Base Node, the MAC_JOIN.indication is passed from the MAC layer to indicate that a multicast group
2290 join was initiated by a Service Node. On a Service Node, it is used to indicate that the Base Node is inviting to
2291 join a multicast group.

2292 Depending on device type, this primitive shall have two variants. The first variant below shall be used in Base
2293 Nodes and the second variant is for Service Nodes:

2294 *MAC_JOIN.indication{ConHandle, EUI-48, Type, Data, DataLength, AE}*

2295 *MAC_JOIN.indication{ConHandle, Type, Data, DataLength, AE}*

2296 The *ConHandle* is a unique identifier interchanged to uniquely identify the multicast group being indicated.
2297 It has a valid meaning only in the MAC SAP, used to have a reference to this connection between different
2298 primitives.

2299 The *EUI-48* parameter indicates which device on the Subnetwork wishes to establish a connection.

2300 The *Type* parameter is an identifier used to define the type of the Convergence layer that should be used for
2301 this request. This parameter is 1 byte long and it is received in the MUL.TYPE field of the connection request.

2302 The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote
2303 CL and the local CL. This parameter is received in the MUL.DATA field of the connection request.

2304 The *DataLength* parameter is the length of the Data parameter in bytes.

2305 The *AE* parameter indicates whether or not the information transmitted in this connection is encrypted. It is
2306 a Boolean type with a value of true indicating that encryption will be used.

2307 **4.5.2.4.5 MAC_JOIN.response**

2308 The MAC_JOIN.response is passed to the MAC layer to respond with a MAC_JOIN.indication. Depending on
2309 device type, this primitive could have either of the two forms given below. The first one shall be used in
2310 Service Node and the second one in Base Node implementations.

2311 The semantics of this primitive are as follows:

2312 *MAC_JOIN.response{ConHandle, Answer, AE}*

2313 *MAC_JOIN.response (ConHandle, EUI, Answer, AE)*

2314 The *ConHandle* parameter is the same as the one that was received in the *MAC_JOIN.indication*.

2315 *EUI* is the EUI-48 of Service Node that requested the multicast group join.

2316 The *Answer* parameter is used to notify the MAC of the action to be taken for this join request. This parameter
2317 may have one of the values depicted below.

2318 The *AE* parameter indicates whether or not the information transmitted in this connection is encrypted. It is
2319 a Boolean type with a value of true indicating that encryption will be used.

2320 **Table 44 - 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.

2321 **4.5.2.5 MAC_LEAVE**

2322 **4.5.2.5.1 General**

2323 The *MAC_LEAVE* primitives are used to leave a broadcast or multicast connection.

2324 **4.5.2.5.2 MAC_LEAVE.request**

2325 The *MAC_LEAVE.request* primitive is used to leave a multicast or broadcast traffic. Depending on device type,
2326 this primitive could have either of the two forms given below. The first one shall be used in Service Node and
2327 the second on in Base Node implementations.

2328 The semantics of this primitive are as follows:

2329 *MAC_LEAVE.request{ConHandle}*

2330 *MAC_LEAVE.request{ConHandle, EUI}*

2331 The *ConHandle* parameter specifies the connection to be left. This handle is the one that was obtained during
2332 the *MAC_JOIN* primitives.

2333 *EUI* is the EUI-48 of Service Node to remove from multicast group.

2334 **4.5.2.5.3 MAC_LEAVE.confirm**

2335 The *MAC_LEAVE.confirm* primitive is received to confirm that the *MAC_LEAVE.request* operation has
2336 finished.

2337 The semantics of this primitive are as follows:

2338 *MAC_LEAVE.confirm{ConHandle, Result}*

2339 The *ConHandle* parameter specifies the connection released. This handle is the one that was obtained during
2340 the *MAC_JOIN* primitives.

2341 The *Result* parameter may have one of the values in Table 45.

2342 **Table 45 - Values of the *Result* parameter in *MAC_LEAVE.confirm* primitive**

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

2343 After the reception of the *MAC_LEAVE.confirm*, the *ConHandle* is no longer valid and should not be used.

2344 **4.5.2.5.4 MAC_LEAVE.indication**

2345 The *MAC_LEAVE.indication* primitive is used to leave a multicast or broadcast traffic. Depending on device
2346 type, this primitive could have either of the two forms given below. The first one shall be used in Service
2347 Node and the second one in Base Node implementations.

2348 The semantics of this primitive are as follows:

2349 *MAC_LEAVE.indication{ConHandle}*

2350 *MAC_LEAVE.indication{ConHandle, EUI}*

2351 The *ConHandle* parameter is the same as that received in *MAC_JOIN.confirm* or *MAC_JOIN.indication*. This
2352 handle is the one that was obtained during the *MAC_JOIN* primitives.

2353 *EUI* is the EUI-48 of Service Node to remove from multicast group.

2354 **4.5.3 Base Node signalling primitives**

2355 **4.5.3.1 General**

2356 This section specifies MAC-SAP primitives that are only available in the Base Node.

2357 **4.5.3.2 MAC_REDIRECT.response**

2358 The *MAC_REDIRECT.response* primitive is used to answer to a *MAC_ESTABLISH.indication* and redirects the
2359 connection from the Base Node to another Service Node on the Subnetwork.

2360 The semantics of this primitive are as follows:

2361 *MAC_REDIRECT.reponse{ConHandle, EUI-48, Data, DataLength}*

2362 The *ConHandle* is the one passed in the MAC_ESTABLISH.indication primitive to which it is replying.
2363 *EUI-48* indicates the Service Node to which this connection establishment should be forwarded. The Base
2364 Node should perform a direct connection setup between the source of the connection establishment and the
2365 Service Node indicated by *EUI-48*.

2366 The *Data* parameter is a general purpose buffer to be interchanged for the negotiation between the remote
2367 CL and the Base Node CL. This parameter is received in the CON.DATA field of the connection request.

2368 The *DataLength* parameter is the length of the *Data* parameter in bytes.

2369 Once this primitive has been used, the *ConHandle* is no longer valid.

2370 **4.5.4 Service and Base Nodes data primitives**

2371 **4.5.4.1 General**

2372 The following subsections describe how a Service Node or Base Node passes data between the Convergence
2373 layer and the MAC layer.

2374 **4.5.4.2 MAC_DATA.request**

2375 The MAC_DATA.request primitive is used to initiate the transmission process of data over a connection.

2376 The semantics of the primitive are as follows:

2377 *MAC_DATA.request*{*ConHandle*, *Data*, *DataLength*, *Priority*, *TimeReference*}

2378 The *ConHandle* parameter specifies the connection to be used for the data transmission. This handle is the
2379 one that was obtained during the connection establishment primitives.

2380 The *Data* parameter is a buffer of octets that contains the CL data to be transmitted through this connection.

2381 The *DataLength* parameter is the length of the *Data* parameter in octets.

2382 *Priority* indicates the priority of the data to be sent when using the CSMA access scheme, i.e. the parameter
2383 only has meaning when the connection was established with *CfBytes* = 0.

2384 The *TimeReference* parameter is the time reference to interchange with the data. This *TimeReference*
2385 parameter is optional; it is possible not sending any time reference. From the primitive point of view the act
2386 of not including a time reference will be considered a *NULL* time reference. The way to interchange this
2387 parameter in the primitive not losing precision and its absolute meaning are specific to the implementation.

2388 **4.5.4.3 MAC_DATA.confirm**

2389 The MAC_DATA.confirm primitive is used to confirm that the transmission process of the data has completed.

2390 The semantics of the primitive are as follows:

2391 *MAC_DATA.confirm*{*ConHandle*, *Data*, *Result*}

2392 The *ConHandle* parameter specifies the connection that was used for the data transmission. This handle is
2393 the one that was obtained during the connection establishment primitives.

2394 The *Data* parameter is a buffer of octets that contains the CL data that where to be transmitted through this
2395 connection.

2396 The *Result* parameter indicates the result of the transmission. This can take one of the values given in Table
2397 42.

2398 **Table 46 - Values of the *Result* parameter in *MAC_DATA.confirm* primitive**

Result	Description
<i>Success</i> = 0	The send was successful.
<i>Timeout</i> = 2	The send process timed out.

2399 **4.5.4.4 MAC_DATA.indication**

2400 The *MAC_DATA.indication* primitive notifies the reception of data through a connection to the CL.

2401 The semantics of the primitive are as follows:

2402 *MAC_DATA.indication*{*ConHandle*, *Data*, *DataLength*, *TimeReference*}

2403 The *ConHandle* parameter specifies the connection where the data was received. This handle is the one that
2404 was obtained during the connection establishment primitives.

2405 The *Data* parameter is a buffer of octets that contains the CL data received through this connection.

2406 The *DataLength* parameter is the length of the *Data* parameter in octets.

2407 The *TimeReference* parameter is the time reference interchanged with the data. This *TimeReference*
2408 parameter is optional; it is possible not receiving any time reference. From the primitive point of view the
2409 act of not indicating a time reference will be considered a *NULL* time reference. The way to interchange this
2410 parameter in the primitive not loosing precision and its absolute meaning are specific to the implementation.

2411 **4.5.5 MAC Layer Management Entity SAPs**

2412 **4.5.5.1 General**

2413 The following primitives are all optional.

2414 The aim is to allow an external management entity to control Registration and Promotion of the Service
2415 Node, demotion and Unregistration of a Service Node. The MAC layer would normally perform this
2416 automatically; however, in some situations/applications it could be advantageous if this could be externally
2417 controlled. Indications are also defined so that an external entity can monitor the status of the MAC.

2418 4.5.5.2 MLME_REGISTER

2419 4.5.5.2.1 General

2420 The MLME_REGISTER primitives are used to perform Registration and to indicate when Registration has been
2421 performed.

2422 4.5.5.2.2 MLME_REGISTER.request

2423 The MLME_REGISTER.request primitive is used to trigger the Registration process to a Subnetwork through
2424 a specific Switch Node. This primitive may be used for enforcing the selection of a specific Switch Node that
2425 may not necessarily be used if the selection is left automatic. The Base Node MLME function does not export
2426 this primitive.

2427 The semantics of the primitive could be either of the following:

2428 *MLME_REGISTER.request{ }*

2429 Invoking this primitive without any parameter simply invokes the Registration process in MAC and leaves the
2430 selection of the Subnetwork and Switch Node to MAC algorithms. Using this primitive enables the MAC to
2431 perform fully automatic Registration if such a mode is implemented in the MAC.

2432 *MLME_REGISTER.request{SNA}*

2433 The *SNA* parameter specifies the Subnetwork to which Registration should be performed. Invoking the
2434 primitive in this format commands the MAC to register only to the specified Subnetwork.

2435 *MLME_REGISTER.request{SID}*

2436 The *SID* parameter is the SID (Switch Identifier) of the Switch Node through which Registration needs to be
2437 performed. Invoking the primitive in this format commands the MAC to register only to the specified Switch
2438 Node.

2439 4.5.5.2.3 MLME_REGISTER.confirm

2440 The MLME_REGISTER.confirm primitive is used to confirm the status of completion of the Registration
2441 process that was initiated by an earlier invocation of the corresponding request primitive. The Base Node
2442 MLME function does not export this primitive.

2443 The semantics of the primitive are as follows:

2444 *MLME_REGISTER.confirm{Result, SNA, SID}*

2445 The *Result* parameter indicates the result of the Registration. This can take one of the values given in Table
2446 43.

2447

Table 47 - Values of the *Result* parameter in MLME_REGISTER.confirm primitive

<i>Result</i>	Description
<i>Done = 0</i>	Registration to specified SNA through specified Switch is completed successfully.
<i>Timeout =2</i>	Registration request timed out .
<i>Rejected=1</i>	Registration request is rejected by Base Node of specified SNA.
<i>NoSNA=8</i>	Specified SNA is not within range.
<i>NoSwitch=9</i>	Switch Node with specified EUI-48 is not within range.

2448 The *SNA* parameter specifies the Subnetwork to which Registration is performed. This parameter is of
2449 significance only if *Result=0*.

2450 The *SID* parameter is the SID (Switch Identifier) of the Switch Node through which Registration is performed.
2451 This parameter is of significance only if *Result=0*.

2452 4.5.5.2.4 MLME_REGISTER.indication

2453 The MLME_REGISTER.indication primitive is used to indicate a status change in the MAC. The Service Node
2454 is now registered to a Subnetwork.

2455 The semantics of the primitive are as follows:

2456
$$MLME_REGISTER.indication\{SNA, SID\}$$

2457 The *SNA* parameter specifies the Subnetwork to which Registration is performed.

2458 The *SID* parameter is the SID (Switch Identifier) of the Switch Node through which Registration is performed.

2459 4.5.5.3 MLME_UNREGISTER

2460 4.5.5.3.1 General

2461 The MLME_UNREGISTER primitives are used to perform deregistration and to indicate when deregistration
2462 has been performed.

2463 4.5.5.3.2 MLME_UNREGISTER.request

2464 The MLME_UNREGISTER.request primitive is used to trigger the Unregistration process. This primitive may
2465 be used by management entities if they require the Node to unregister for some reason (e.g. register through
2466 another Switch Node or to another Subnetwork). The Base Node MLME function does not export this
2467 primitive.

2468 The semantics of the primitive are as follows:

2469
$$MLME_UNREGISTER.request\{\}$$

2470 **4.5.5.3.3 MLME_UNREGISTER.confirm**

2471 The MLME_UNREGISTER.confirm primitive is used to confirm the status of completion of the unregister
2472 process initiated by an earlier invocation of the corresponding request primitive. The Base Node MLME
2473 function does not export this primitive.

2474 The semantics of the primitive are as follows:

2475 *MLME_UNREGISTER.confirm{Result}*

2476 The *Result* parameter indicates the result of the Registration. This can take one of the values given in Table
2477 44.

2478 **Table 48 - Values of the *Result* parameter in MLME_UNREGISTER.confirm primitive**

Result	Description
<i>Done = 0</i>	Unregister process completed successfully.
<i>Timeout =2</i>	Unregister process timed out .
<i>Redundant=10</i>	The Node is already in <i>Disconnected</i> functional state and does not need to unregister.

2479 On generation of MLME_UNREGISTER.confirm, the MAC layer shall not perform any automatic actions that
2480 may invoke the Registration process again. In such cases, it is up to the management entity to restart the
2481 MAC functionality with appropriate MLME_REGISTER primitives.

2482 **4.5.5.3.4 MLME_UNREGISTER.indication**

2483 The MLME_UNREGISTER.indication primitive is used to indicate a status change in the MAC. The Service Node
2484 is no longer registered to a Subnetwork.

2485 The semantics of the primitive are as follows:

2486 *MLME_UNREGISTER.indication{}*

2487 **4.5.5.4 MLME_PROMOTE**

2488 **4.5.5.4.1 General**

2489 The MLME_PROMOTE primitives are used to perform promotion and to indicate when promotion has been
2490 performed.

2491 **4.5.5.4.2 MLME_PROMOTE.request**

2492 The MLME_PROMOTE.request primitive is used to trigger the promotion process in a Service Node that is in
2493 a *Terminal* functional state. Implementations may use such triggered promotions to optimize Subnetwork
2494 topology from time to time. The value of PRO.PNA in the promotion message sent to the Base Node is
2495 undefined and implementation-specific.

2496 The MLME_PROMOTE.request primitive can also be used from a node that is already in a *Switch* state to ask
2497 the BN for a Beacon PDU modulation change.

2498 Base Node can use this primitive to ask a node to change its state from *Terminal* to *Switch* or, if the node is
2499 already in the *Switch* state, to adopt a new Beacon PDU modulation scheme.

2500 The semantics of the primitive can be either of the following:

2501 $MLME_PROMOTE.request\{\}$

2502 $MLME_PROMOTE.request\{BCN_MODE\}$

2503 $MLME_PROMOTE.request\{EUI-48, BCN_MODE\}$

2504 The EUI-48 parameter shall be used only by the Base Node to specify the address of the Node to which this
2505 promotion request shall be addressed. The MAC shall internally transfer this to an address used by the MAC
2506 layer.

2507 The BCN_MODE parameter specifies the Beacon PDU modulation scheme. If the primitive is called by a node
2508 in *Switch* state, this parameter indicates the requested Beacon PDU modulation scheme from the *Switch*
2509 node to the Base Node. If the primitive is called by the Base Node, this parameter indicates the modulation
2510 scheme that shall be communicated to the node during the promotion process or during the Beacon PDU
2511 modulation change process.

2512 Allowed values for BCN_MODE parameter are listed in Table 49.

2513 **Table 49 - Values of the BCN_MODE parameter in MLME_PROMOTE.request primitive.**

BCN_MODE	Description
DBPSK_F = 0	BCN will be sent using DBPSK modulation with convolutional encoding enabled and robust mode disabled.
R_DBPSK = 1	BCN will be sent using DBPSK modulation with robust mode enabled.
R_DQPSK = 2	BCN will be sent using DQPSK modulation with robust mode enabled.

2514 **4.5.5.4.3 MLME_PROMOTE.confirm**

2515 The MLME_PROMOTE.confirm primitive is used to confirm the status of completion of a promotion process
2516 that was initiated by an earlier invocation of the corresponding request primitive.

2517 The semantics of the primitive are as follows:

2518 $MLME_PROMOTE.confirm\{Result\}$

2519 The *Result* parameter indicates the result of the Registration. This can take one of the values given in Table
2520 45.

2521 **Table 50 - Values of the Result parameter in MLME_PROMOTE.confirm primitive**

Result	Description
<i>Done = 0</i>	Node is promoted to Switch function successfully.
<i>Timeout =1</i>	Promotion process timed out.
<i>Rejected=2</i>	The Base Node rejected promotion request.
<i>No Such Device = 4</i>	A device with the destination address cannot be found.
<i>Redundant=10</i>	This device is already functioning as Switch Node.
<i>OutofRange=12</i>	Specified BCN_MODE is out of acceptable range.

2522 In case an already promoted switch, which is requesting a Beacon PDU modulation change, receives an
2523 MLME_PROMOTE.confirm{} rejecting the request, only the change request is supposed to be rejected, so the
2524 node shall continue sending the Beacon PDU as previously.

2525 **4.5.5.4.4 MLME_PROMOTE.indication**

2526 The MLME_PROMOTE.indication primitive is used to indicate a status change in the MAC. The Service Node
2527 is now operating as a Switch. This primitive is not generated if a Beacon PDU modulation change occurs.

2528 The semantics of the primitive are as follows:

2529 *MLME_PROMOTE.indication{}*

2530 **4.5.5.5 MLME_DEMOTE**

2531 **4.5.5.5.1 General**

2532 The MLME_DEMOTE primitives are used to perform demotion and to indicate when demotion has been
2533 performed.

2534 **4.5.5.5.2 MLME_DEMOTE.request**

2535 The MLME_DEMOTE.request primitive is used to trigger a demotion process in a Service Node that is in a
2536 *Switch* functional state. This primitive may be used by management entities to enforce demotion in cases
2537 where the Node's default functionality does not automatically perform the process.

2538 The semantics of the primitive are as follows:

2539 *MLME_DEMOTE.request{}*

2540 **4.5.5.5.3 MLME_DEMOTE.confirm**

2541 The MLME_DEMOTE.confirm primitive is used to confirm the status of completion of a demotion process
2542 that was initiated by an earlier invocation of the corresponding request primitive.

2543 The semantics of the primitive are as follows:

2544
$$MLME_DEMOTE.confirm\{Result\}$$

2545 The *Result* parameter indicates the result of the demotion. This can take one of the values given in Table 46.

2546 **Table 51 - Values of the *Result* parameter in MLME_DEMOTE.confirm primitive**

Result	Description
<i>Done = 0</i>	Node is demoted to Terminal function successfully.
<i>Timeout =1</i>	Demotion process timed out.
<i>Redundant=10</i>	This device is already functioning as Terminal Node.

2547 When a demotion has been triggered using the MLME_DEMOTE.request, the Terminal will remain demoted.

2548 **4.5.5.5.4 MLME_DEMOTE.indication**

2549 The MLME_DEMOTE.indication primitive is used to indicate a status change in the MAC. The Service Node is
2550 now operating as a Terminal.

2551 The semantics of the primitive are as follows:

2552
$$MLME_DEMOTE.indication\{\}$$

2553 **4.5.5.6 MLME_RESET**

2554 **4.5.5.6.1 General**

2555 The MLME_RESET primitives are used to reset the MAC into a known good status.

2556 **4.5.5.6.2 MLME_RESET.request**

2557 The MLME_RESET.request primitive results in the flushing of all transmit and receive buffers and the resetting
2558 of all state variables. As a result of invoking of this primitive, a Service Node will transit from its present
2559 functional state to the *Disconnected* functional state.

2560 The semantics of the primitive are as follows:

2561
$$MLME_RESET.request\{\}$$

2562 **4.5.5.6.3 MLME_RESET.confirm**

2563 The MLME_RESET.confirm primitive is used to confirm the status of completion of a reset process that was
2564 initiated by an earlier invocation of the corresponding request primitive. On the successful completion of the
2565 reset process, the MAC entity shall restart all functions starting from the search for a Subnetwork (4.3.1).

2566 The semantics of the primitive are as follows:

2567 *MLME_RESET.confirm{Result}*

2568 The *Result* parameter indicates the result of the reset. This can take one of the values given below.

2569 **Table 52 - Values of the *Result* parameter in MLME_RESET.confirm primitive**

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

2570 **4.5.5.7 MLME_GET**

2571 **4.5.5.7.1 General**

2572 The MLME_GET primitives are used to retrieve individual values from the MAC, such as statistics.

2573 **4.5.5.7.2 MLME_GET.request**

2574 The MLME_GET.request queries information about a given PIB attribute.

2575 The semantics of the primitive are as follows:

2576 *MLME_GET.request{PIBAttribute}*

2577 The *PIBAttribute* parameter identifies specific attributes as listed in the *Id* fields of tables that list PIB
2578 attributes (Section 6.2.3).

2579 **4.5.5.7.3 MLME_GET.confirm**

2580 The MLME_GET.confirm primitive is generated in response to the corresponding MLME_GET.request
2581 primitive.

2582 The semantics of this primitive are as follows:

2583 *MLME_GET.confirm{status, PIBAttribute, PIBAttributeValue}*

2584 The *status* parameter reports the result of requested information and can have one of the values given in
2585 Table 48.

2586

Table 53 - Values of the *status* parameter in MLME_GET.confirm primitive

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

2587 The *PIBAttribute* parameter identifies specific attributes as listed in *Id* fields of tables that list PIB attributes
2588 (Section 6.2.3.5).

2589 The *PIBAttributeValue* parameter specifies the value associated with a given *PIBAttribute*.

2590 4.5.5.8 MLME_LIST_GET

2591 4.5.5.8.1 General

2592 The MLME_LIST_GET primitives are used to retrieve a list of values from the MAC.

2593 4.5.5.8.2 MLME_LIST_GET.request

2594 The MLME_LIST_GET.request queries for a list of values pertaining to a specific class. These special classes of
2595 PIB attributes are listed in Table 100.

2596 The semantics of the primitive are as follows:

2597 *MLME_LIST_GET.request{PIBListAttribute}*

2598 The *PIBListAttribute* parameter identifies a specific list that is requested by the management entity. The
2599 possible values of *PIBListAttribute* are listed in 6.2.3.5.

2600 4.5.5.8.3 MLME_LIST_GET.confirm

2601 The MLME_LIST_GET.confirm primitive is generated in response to the corresponding
2602 MLME_LIST_GET.request primitive.

2603 The semantics of this primitive are as follows:

2604 *MLME_LIST_GET.confirm{status, PIBListAttribute, PIBListAttributeValue}*

2605 The *status* parameter reports the result of requested information and can have one of the values given in
2606 Table 49

2607 Table 54 - Values of the *status* parameter in MLME_LIST_GET.confirm primitive

Result	Description
<i>Done = 0</i>	Parameter read successfully.
<i>Failed =1</i>	Parameter read failed due to internal implementation reasons.

Result	Description
<i>BadAttr=11</i>	Specified <i>PIBListAttribute</i> is not supported.

2608 The *PIBListAttribute* parameter identifies a specific list as listed in the *Id* field of Table 100.

2609 The *PIBListAttributeValue* parameter contains the actual listing associated with a given *PIBListAttribute*

2610 **4.5.5.9 MLME_SET**

2611 **4.5.5.9.1 General**

2612 The MLME_SET primitives are used to set configuration values in the MAC.

2613 **4.5.5.9.2 MLME_SET.request**

2614 The MLME_SET.requests information about a given PIB attribute.

2615 The semantics of the primitive are as follows:

2616 *MLME_SET.request{PIBAttribute, PIBAttributeValue}*

2617 The *PIBAttribute* parameter identifies a specific attribute as listed in the *Id* fields of tables that list PIB
2618 attributes (Section 6.2.3).

2619 The *PIBAttributeValue* parameter specifies the value associated with given *PIBAttribute*.

2620 **4.5.5.9.3 MLME_SET.confirm**

2621 The MLME_SET.confirm primitive is generated in response to the corresponding MLME_SET.request
2622 primitive.

2623 The semantics of this primitive are as follows:

2624 *MLME_SET.confirm{result}*

2625 The *result* parameter reports the result of requested information and can have one of the values given in
2626 Table 50.

2627 **Table 55 - Values of the *Result* parameter in MLME_SET.confirm primitive**

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

2628 The *PIBAttribute* parameter identifies a specific attribute as listed in the *Id* fields of tables that list PIB
2629 attributes (Section 6.2.3).

2630 The *PIBAttributeValue* parameter specifies the value associated with a given *PIBAttribute*.

2631 **4.6 MAC procedures**

2632 **4.6.1 Registration process**

2633 **4.6.1.1 General**

2634 The initial Service Node start-up (4.3.1) is followed by a Registration process. A Service Node in a
2635 *Disconnected* functional state shall transmit a REG control packet to the Base Node in order to get itself
2636 included in the Subnetwork. Since no LNID or SID is allocated to a Service Node at this stage, the PKT.LNID
2637 field shall be set to all 1s and the PKT.SID field shall contain the SID of the Switch Node through which it seeks
2638 attachment to the Subnetwork.

2639 Base Nodes may use a Registration request as an authentication mechanism. However this specification does
2640 not recommend or forbid any specific authentication mechanism and leaves this choice to implementations.

2641 For all successfully accepted Registration requests, the Base Node shall allocate an LNID that is unique within
2642 the domain of the Switch Node through which the attachment is realized. This LNID shall be indicated in the
2643 PKT.LNID field of response (REG_RSP). The assigned LNID, in combination with the SID of the Switch Node
2644 through which the Service Node is registered, would form the NID of the registering Node.

2645 Registration is a three-way process. The Base Node answers to the REG_REQ - registration request - sent by
2646 a Service Node by means of a REG_RSP message, which shall be acknowledged by the Service Node with a
2647 REG_ACK message.

2648 Service Nodes report their capabilities to the Base Node during registration (REG_REQ), as specified in
2649 4.4.2.6.3. On top of that, a Base Node is able to configure some parameters in Service Nodes when answering
2650 (REG_RSP) to a registration request.

- 2651 • Dynamic robustness-management is enabled by default. Nonetheless, the Base Node may disable
2652 dynamic robustness-management and fix a specific modulation scheme, thus not allowing Service
2653 Node(s) to dynamically switch to a different modulation scheme.

2654 The configured value is stored by the Service Node as "*macRobustnessManagement*".

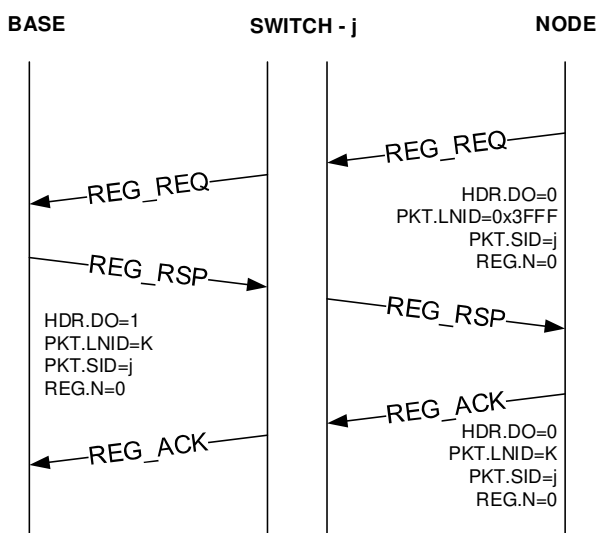
- 2655 • *Segmentation And Reassembly (SAR)* packet size: The packet size used by Convergence Layer's SAR
2656 Service is not configured by the Base Node by default. Nonetheless, a Base Node may fix a specific
2657 SAR packet size if required. For more information about Convergence Layer's SAR Service, please
2658 refer to section 5.6.2.1.3

2660 The configured value is stored by the Service Node as "*macSARsize*".

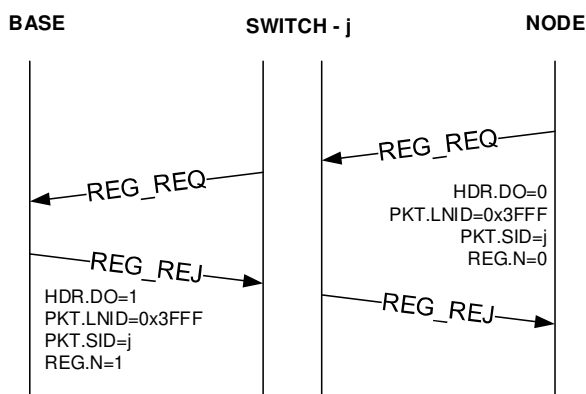
2662 Configuration of the two parameters mentioned above during registration provides a static network
 2663 configuration. This configuration can be changed by the Base Node, either starting a new registration process
 2664 or setting the corresponding Service Node's PIB variables remotely.

2665 Figure 71 represents a successful Registration process and Figure 72 shows a Registration request that is
 2666 rejected by the Base Node. Details on specific fields that distinguish one Registration message from the other
 2667 are given in Table 20. Figure 71 also denotes the security-related steps that pertain to the registration
 2668 process. The registration process security-related steps are explained in Section 4.6.1.2.

2669 The REG control packet, in all its usage variants, is transmitted unencrypted, but specified fields (REG.SWK
 2670 and REG.WK) are encrypted with context-specific encryption keys as explained in Section 4.4.2.6.3. The
 2671 encryption of REG.WK in REG_RSP, its decryption at the receiving end and subsequent encrypted
 2672 retransmission using a different encryption key authenticates that the REG_ACK is from the intended
 2673 destination.



2674
 2675 **Figure 71 - Registration process accepted**



2676
 2677 **Figure 72 - Registration process rejected**

2678 When assigning an LNID, the Base Node shall not reuse an LNID released by an unregister process before
 2679 (*macCtrlMsgFailTime + macMinCtlReTxTimer*) seconds, to ensure that all retransmitted packets have left the
 2680 Subnetwork. Similarly, the Base Node shall not reuse an LNID released by the Keep-Alive process before

2681 $T_{\text{keep_alive}}$ seconds, using the last known acknowledged $T_{\text{keep_alive}}$ value, or if larger, the last unacknowledged
2682 $T_{\text{keep_alive}}$, for the Service Node using the LNID. When security is being used in the network, the Base Node
2683 shall not reuse a LNID without first changing the Subnetwork Working Key.

2684 During network startup where the whole network is powered on at once, there will be considerable
2685 contention for the medium. It is recommended to add randomness to the first REG_REQ transmission, as
2686 well as to all subsequent retransmissions. It is recommended to wait a random delay before the first
2687 REG_REQ message. This delay should be in range from 0 to at least 10% of *macCtrlMsgFailTime*. Similarly a
2688 random delay may be added to each retransmission.

2689 **4.6.1.2 Security registration process**

2690 Figure 71 represents the registration process. When security profile 1 or 2 is utilized, additional action is
2691 required by the Base and Terminal Nodes to ensure successful registration.

- 2692 1. The Terminal Node generates a challenge (see Section 0)
- 2693 2. The challenge is included in the REG_REQ and the REG_REQ is authenticated with REGK.
- 2694 3. The Base Node validates that REG_REQ is properly authenticated.
- 2695 4. The SWK and WK are key wrapped with KWK. The REG_RSP is authenticated with REGK and the
2696 Terminal Node challenge is concatenated.
- 2697 5. The Terminal Node validates that REG_RSP is properly authenticated, including the concatenated
2698 challenge.
- 2699 6. The Terminal Node updates WK and SWK.
- 2700 7. The REG_ACK is authenticated with WK. The first Nonce is required for AES-CCM (Set to 0, then
2701 counted up for every packet.)
- 2702 8. The Base Node validates REG_ACK. The registration is invalidated on error

2703 **4.6.2 Unregistration process**

2704 At any point in time, either the Base Node or the Service Node may decide to close an existing registration.
2705 This version of the specification does not provide provision for rejecting an unregistration request. The
2706 Service Node or Base Node that receives an unregistration request shall acknowledge its receipt and take
2707 appropriate actions.

2708 Following a successful unregistration, a Service Node shall move back from its present functional state to a
2709 *Disconnected* functional state and the Base Node may re-use any resources that were reserved for the
2710 unregistered Node.

2711 Figure 73 shows a successful unregistration process initiated by a Service Node and Figure 74 shows an
2712 unregistration process initiated by the Base Node. Details on specific fields that identify unregistration
2713 requests in REG control packets are given in Table 21.




Figure 73 - Unregistration process initiated by a Terminal Node

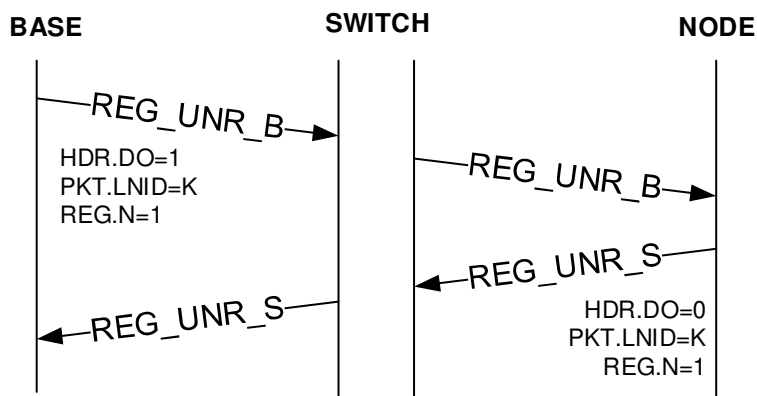


Figure 74 - Unregistration process initiated by the Base Node

4.6.3 Promotion process

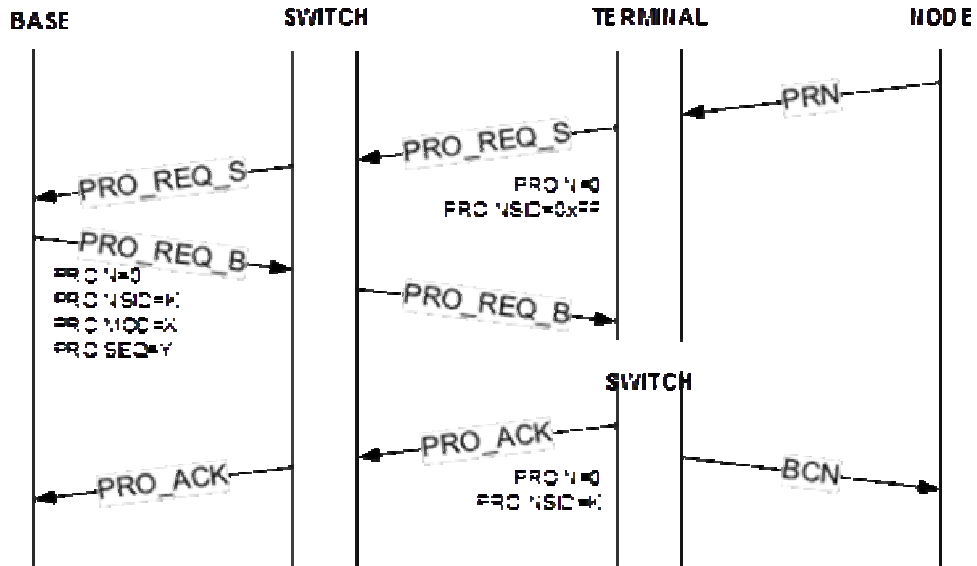
A Service Node that does not receive any BPDUs may transmit PNPDU. Any Terminal Node receiving PNPDU may generate a promotion request towards Base Node, which upon acceptance from Base Node, will result in transition of the requesting Terminal Node to Switch and therefore scale the Subnetwork to facilitate PNPDU transmitting Service Node to join.

Note: A Subnetwork that operates in backward compatibility-mode as enumerated in 4.8, shall silently discard PNPDU that indicate lack of support for backward compatibility-mode i.e. PNH.VER = 1 and PNH.CAP_BC = 0.

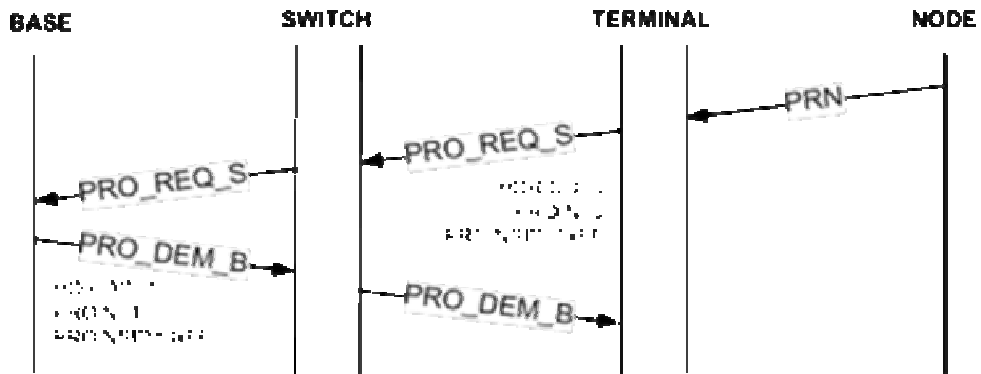
The Base Node examines promotion requests during a period of time. It may use the address of the new Terminal, provided in the promotion-request packet, to decide whether or not to accept the promotion. It decides which Service Node shall be promoted, if any, sending a promotion response. The other Nodes do not receive any response to their 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 1s. 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.

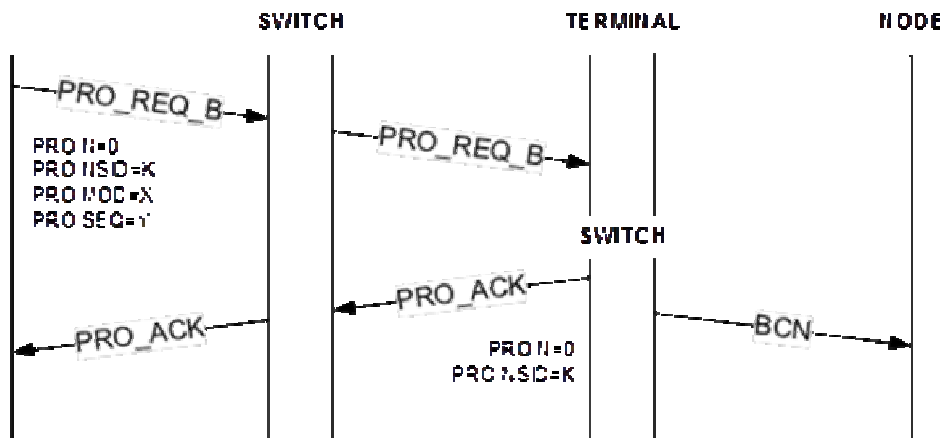
2736 When reusing LSIDs that have been released by a demotion process, the Base Node shall not allocate the
 2737 LSID until after $(macCtrlMsgFailTime + macMinCtrlReTxTimer)$ seconds to ensure all retransmit packets that
 2738 might use that LSID have left the Subnetwork.




2739
 2740 Figure 75 - Promotion process initiated by a Service Node



2741
 2742 Figure 76 - Promotion process rejected by the Base Node



2743
 2744 Figure 77 - Promotion process initiated by the Base Node




2745
2746

Figure 78 - Promotion process rejected by a Service Node

2747 **4.6.3.1 BPDU modulation change**

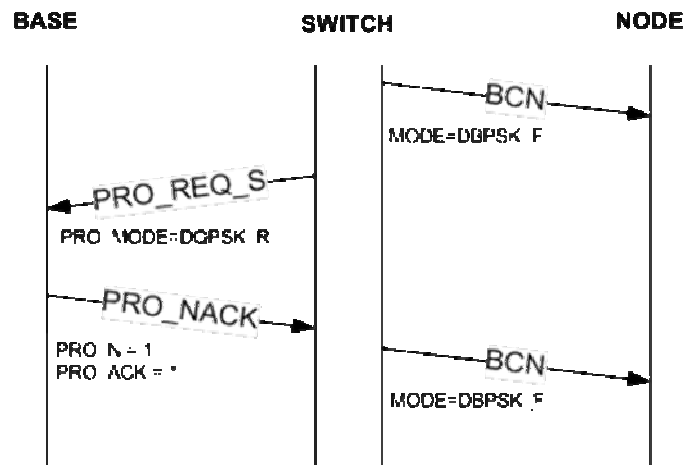
2748 It is possible, for the Switch Node, to change modulation scheme used to send BPDUs. In order to do so a
 2749 new PRO_REQ_S message is sent with an indication of the new desired modulation scheme to be used in
 2750 PRO.MOD field. On reception of this PRO_REQ_S the Base Node shall send a PRO_ACK packet to accept the
 2751 change request, or send a PRO_NACK packet to reject the change request. The Base Node can not indicate a
 2752 new BPDU modulation scheme in the PRO.MOD field that is different from the requested one. In such a case
 2753 the Base Node can accept the requested modulation and initiate another beacon modulation change by itself.
 2754 The Switch would then either acknowledge the reception by sending a PRO_ACK (accept the new modulation)
 2755 or a PRO_NACK packet (reject the new modulation). In case an explicit denial is issued by the Base Node, the
 2756 Switch shall keep on sending the Beacon PDUs without changing to the new modulation scheme. Switch Node
 2757 shall not sent PRO_ACK packet in case of explicit reject.

2758 The Beacon PDU modulation change process can also be initialized by the Base Node. In this case the Switch
 2759 Node shall send a PRO_ACK packet if it can perform the Beacon PDU modulation change otherwise it shall
 2760 send a PRO_NACK.



2761
2762


Figure 79 - BCN modulation change request initiated by the Switch.



2763

2764


Figure 80 - BCN modulation change request initiated by the Switch and rejected by the Base Node.



2765

2766

Figure 81 - BCN modulation change request initiated by the Base Node.



2767

2768

Figure 82 - BCN modulation change request initiated by the Base Node and rejected by the Switch

2769 During a Promotion procedure the Base Node assigns resources to new Switch Node in order to transmit its
2770 BPDUs. These changes shall take effect only on a super-frame boundary. In case these changes require a
2771 change in frame structure, the Base Node shall send a FRA packet to inform the entire network.

2772 **4.6.3.2 Double switching procedure**

2773 Certain Subnetworks may have a mix of device-types between ones that can support Type A PHY frames only
2774 and ones that support both Type A and Type B PHY frames. In such cases, a Switch Node that acts as switching
2775 point for both kinds of devices, may need to transmit BPDUs using both types of PHY frames.

2776 In order to be able to transmit BPDUs using both types of PHY frames, a Switch Node needs to undergo a
2777 second promotion procedure. The first promotion is carried out in the usual manner as enumerated 4.6.3.
2778 When a Switch Node identifies need to transmit its BPDUs in additional modulation scheme, it starts a second
2779 promotion procedure. The Switch Node uses PRO packets with the PRO.DS bit set to one. Additionally, the
2780 PRO_REQ_S packet shall fill LSID of the requesting Switch Node in PRO.SID field.

2781 When a Switch Node has two BPDUs to send it may ask to change modulation scheme only for the robust
2782 beacon, passing from the DBPSK_R to DQPSK_R or viceversa following procedure enumerated in 4.6.3.1.

2783 To stop sending one type of BPDUs, the demotion procedure is used. In this case the PRO.MOD field indicates
2784 which beacon shall not be transmit and the PRO.DS field set to 1 idicates that the node is asking to stop
2785 sending one type of beacon. If the PRO.DS field is set to 0 the node is asking for a full demotion, stop sending
2786 both BPDUs and transit back to *Terminal* funcional state.

2787 By having possibility to provide connectivity for Type A only devices and for devices that require Type B
2788 frames, the Switch Node shall guarantee delivery of multicast and broadcast packets. In simplified
2789 implementations, the Switch Node can transmit these types of packets twice, one with the Type A frame and
2790 one with the Type B. Multicast data shall be switched in conformance to procedures enumerated in 4.6.7.4.3.

2791 For broacast data, the Switch Node shall start to send packets twice after it succesfully performs the double
2792 beacon slot allocation, and it shall stop sending one of the two type of packets when the demotion procedure
2793 is completed for that specific type of frame.

2794 Devices that are able to understand both frames, Type A or Type B may receive same data twice, at each
2795 modulation scheme. Such devices shall be intelligent enough to discard the duplicate receipt of data..

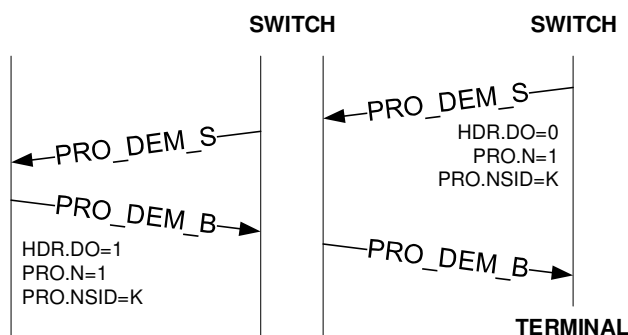
2796 **4.6.4 Demotion process**

2797 The Base Node or a Switch Node may decide to discontinue a switching function at anytime. The demotion
2798 process provides for such a mechanism. The PRO control packet is used for all demotion transactions.

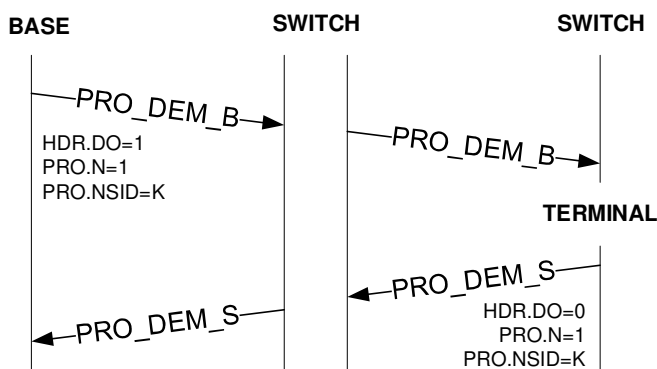
2799 The PRO.NSID field shall contain the SID of the Switch Node that is being demoted as part of the demotion
2800 transaction. The PRO.PNA field is not used in any demotion process transaction and its contents are not
2801 interpreted at either end. PRO.MOD field is not used, it shall be set to zero and not interpreted at either end.

2802 Following successful completion of a demotion process, a Switch Node shall immediately stop the
2803 transmission of beacons and change from a *Switch* funcional state to a *Terminal* funcional state.

2804 The present version of this specification does not specify any explicit message to reject a demotion requested
 2805 by a peer at the other end.



2806
 2807 **Figure 83 - Demotion process initiated by a Service Node**



2808
 2809 **Figure 84 - Demotion process initiated by the Base Node**

2810 4.6.5 Keep-Alive process

2811 4.6.5.1 General

2812 The Keep-Alive process is used to perform two operations:

- 2813 • To detect when a Service Node has left the Subnetwork because of changes to the network
 2814 configuration or because of fatal errors it cannot recover from.
- 2815 • To perform robustness management on each hop in the path to the Service Node.

2816 Service Node shall use one timer, $T_{\text{keep-alive}}$, to detect if it is no longer part of the Subnetwork. If $T_{\text{keep-alive}}$
 2817 expires, the Service Node assumes it has been unregistered by the Base Node and shall enter in the
 2818 *Disconnected* functional state. The timer is started when the Service Node receives the REG_RSP packet with
 2819 value encoded in the REG.TIME field.

2820 The timer is refreshed when any of the following packets has been received with the TIME information
 2821 provided in those packets:

- 2822 • REG_RSP packet (repetition).
- 2823 • ALV_REQ_B packet.
- 2824 • PRO_REQ packet.

2825 The timer is also restarted with the last time received in one of the above packets according to the following
2826 rules:

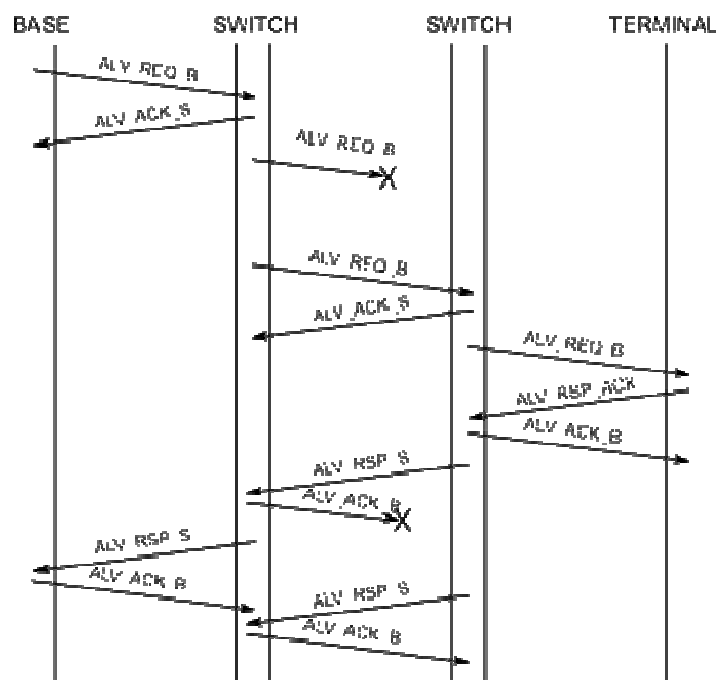
- 2827
- For nodes in Terminal state, the timer is restarted on reception of
 - data on an ARQ connection, which fulfills the following conditions: is originated from the Base Node, is addressed to the node itself and has not yet been acknowledged. Repetitions of the same packet shall not update the timer.
 - a CON_REQ_B, CON_CLS_B, MUL_JOIN_B, MUL_LEAVE_B or SEC_REQ control packet which is addressed to the node itself.

2833 Intermediate Switch nodes restart their timer when transmitting an ALV_RSP_S from an ALV procedure of a
2834 node below them. The timer is restarted with the last time information received in a REG_RSP, ALV_REQ_B
2835 or PRO_REQ packet addressed to the switch node itself.

2836 Each switch along the path to a node keeps track of the switches that are being promoted below it as
2837 described in section 4.3.4.3

2838 The keep alive process has a link level acknowledge, each switch in the path to the target Service Node is
2839 responsible for the retransmissions with the next node, up to *macALVHopRepetitions* retransmissions
2840 (*macALVHopRepetitions* + 1 packets sent). Each retransmission shall be performed in a time equal to a frame
2841 time. On a reception of an ALV_REQ_B/ALV_RSP_S the receiving node shall respond with an
2842 ALV_ACK_S/ALV_ACK_B as soon as possible and with a priority of 0. These retransmissions shall be used to
2843 perform robustness management according to section 7.

2844 If the Service Node identifies that the received ALV_REQ_B/ALV_RSP_S is a retransmit of an already received
2845 packet, the node shall send the related ACK but shall not switch the Alive to the next hop (since it already did
2846 it). The algorithm to detect this situation is up to the manufacturer, as a guideline, it could store the last ALV
2847 operation's data and check if it matches.

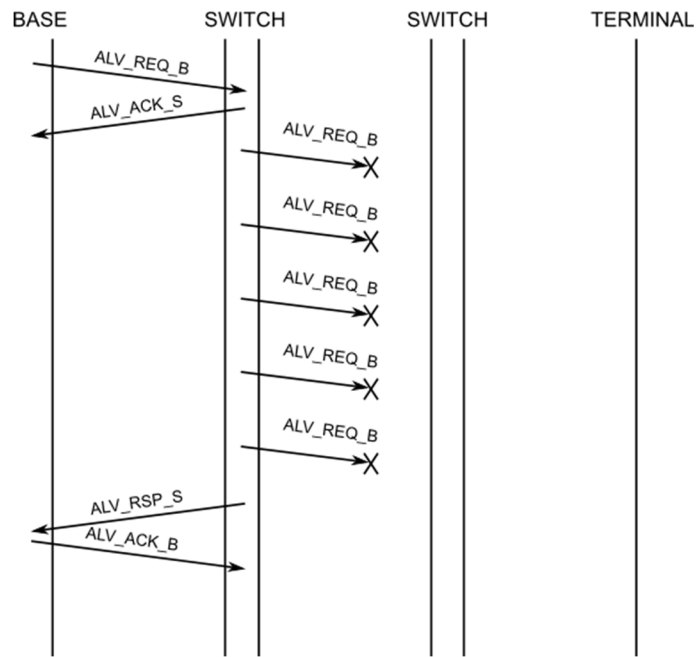


2848

2849

Figure 85 - Successful ALV procedure


2850 If the retransmissions reach the maximum during ALV_REQ_B process, the switch node shall start the
 2851 ALV_RSP_S procedure.



2852

2853

Figure 86 - Failed ALV procedure



2854

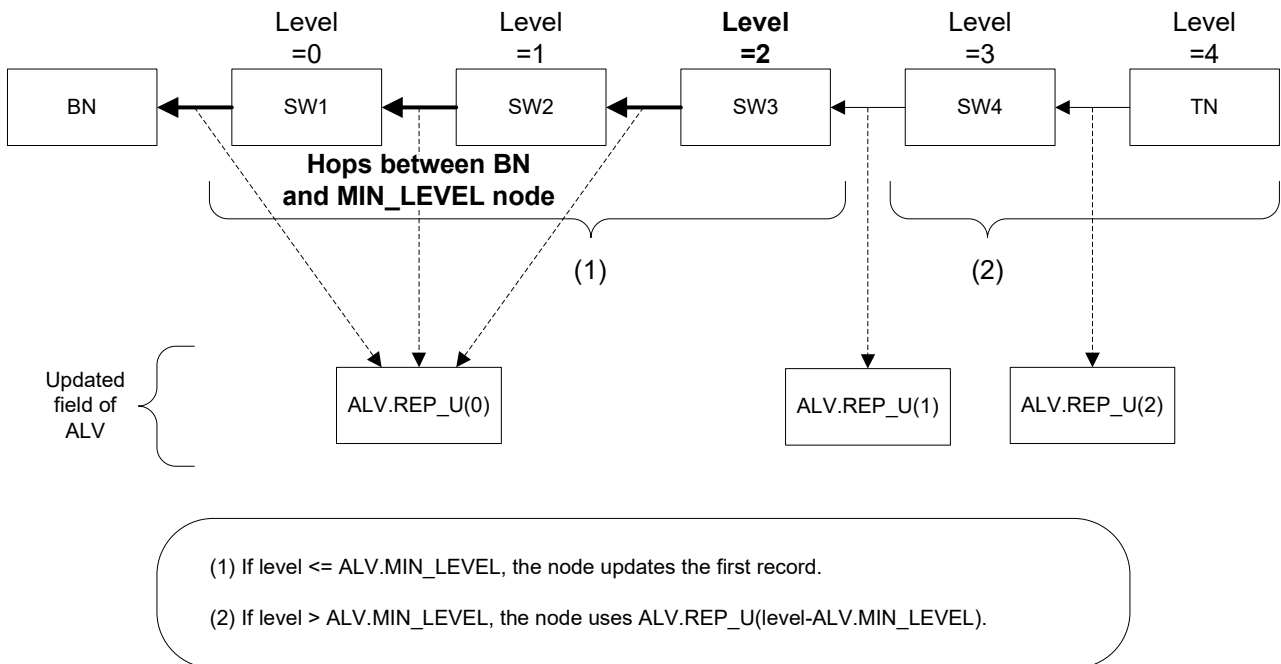
2855

Figure 87 - Failed ALV procedure (uplink)

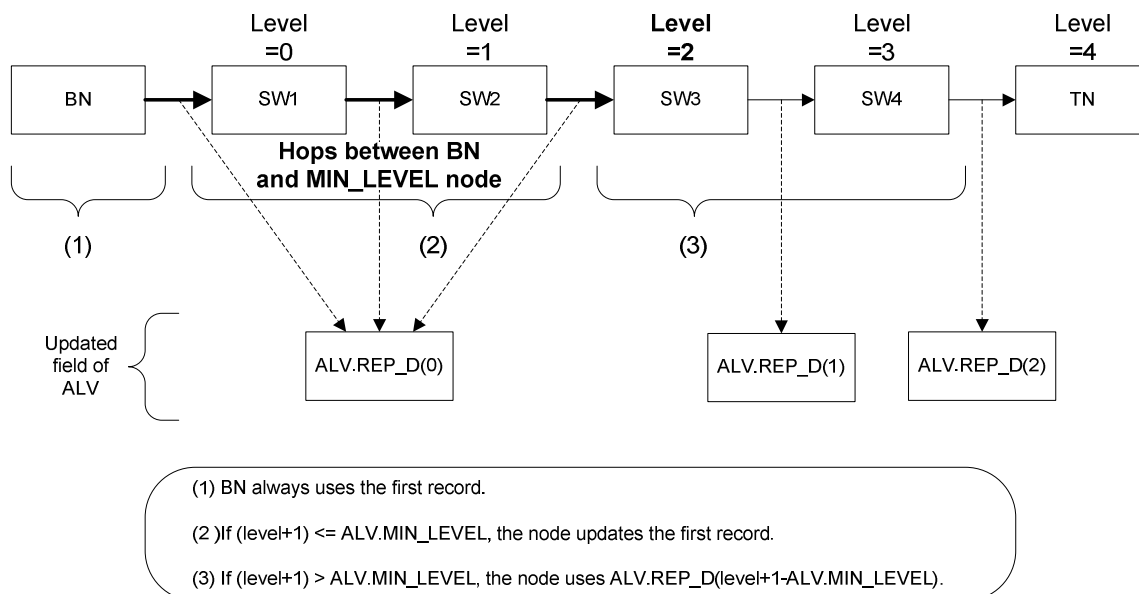
2856 Every time a switch performs a retransmission, it shall decrease the ALV.RTL field before sending the packet,
2857 and fill the record of local retransmissions accordingly. When ALV.RTL reaches 0 and the switch has to
2858 perform a retransmission, it shall discard the packet.

2859 Every switch shall add the information of the retransmissions needed to reach the node in the ALV.REP_D
2860 and ALV.REP_U fields, filling the array with the retransmissions needed for each link, both downlink and
2861 uplink. The Base Node shall form the ALV message with all the registries and the Service Nodes shall fill it
2862 with the values. The rule to know in which record a switch shall add its repetitions is the following.

- 2863 • Uplink
 - 2864 ○ Base node always uses the first record (record number 0: ALV.REP_D(0)).
 - 2865 ○ If the switch's level+1 is equal or lower than ALV.MIN_LEVEL it shall sum its repetitions to
2866 the first record (record number 0: ALV.REP_D(0)/ALV.REP_U(0)).
 - 2867 ○ If the switch's level+1 is greater than ALV.MIN_LEVEL it shall use the records sequentially
2868 from the beginning. A formula to compute which record shall be used: subtract
2869 ALV.MIN_LEVEL from its own level+1. E.g.: If ALV.MIN_LEVEL = 2 and the switch is at level 3
2870 it shall use the record ALV.REP_D(2).
- 2871 • Downlink
 - 2872 ○ If the switch's (or terminal's) level is equal or lower than ALV.MIN_LEVEL it shall sum its
2873 repetitions to the first record.
 - 2874 ○ If the switch's (or terminal's) level is greater than ALV.MIN_LEVEL it shall use the records
2875 sequentially from the beginning. A formula to compute which record shall be used: subtract
2876 ALV.MIN_LEVEL from its own level. E.g.: If ALV.MIN_LEVEL = 2 and the switch is at level 3 it
2877 shall use the record ALV.REP_U(1).



2878
2879 **Figure 88 - ALV.REP_U Example (ALV.MIN_LEVEL = 2)**



2880

2881

Figure 89 - ALV.REP_D Example (ALV.MIN_LEVEL = 2)

2882 The Base Node shall provide the uplink information whenever available in the ALV.REP_U(*) fields using the
 2883 ALV_REQ_B message, this provides connection quality information to the service node. If the retransmission
 2884 at any hop is not available at the time the ALV_REQ_B is sent, the Base Node shall set the ALV.VALU(*) value
 2885 to 0 for that hop, and the Service Node shall ignore this record. The first ALV procedure for a hop after
 2886 registration of a Service Node shall be mark as invalid.

2887 At the end of the process the Base Node receives the ALV_RSP_S of the last switch with information of the
 2888 connectivity of the node and all the hops in its path. This operation is more robust than round-trip control
 2889 packet transaction (CON, REG), so the base node can decide that the node does not have enough connectivity
 2890 and start an unregistration process with it.

2891 The algorithm used by the Base Node to determine when to send ALV_REQ_B messages to registered Service
 2892 Nodes and how to determine the value ALV.TIME, PRO.TIME and REG.TIME is left to implementers.

2893 A Switch Node is required to be able to queue *MACConcurrentAliveProcedure* of each ALV_REQ_B and
 2894 ALV_RSP_S messages at a time. The base node is shall space the ALV_REQ_B queries appropriately.

2895 **4.6.5.2 Use of legacy PRIME 1.3.6 keep-alive mechanism**

2896 In some particular network configurations the use of legacy PRIME 1.3.6 Keep-Alive process instead of the
 2897 PRIME 1.4 Keep-Alive process can be required. For this reason, all PRIME 1.4 implementations must be able
 2898 to implement support for REG.ALV_F field (Section 4.4.2.6.3) in REG control packet. This implies that the Base
 2899 Node shall have the ability to move the entire Subnetwork to ALV procedure listed in Section K.2.5. The
 2900 Subnetwork Keep-Alive process in use can be determined by reading the Base Node PIB attribute
 2901 *macAliveTimeMode*. The configuration of the Keep-Alive process is left to Base Node implementations.


2902 **4.6.6 Connection establishment**

2903 Connection establishment works end-to-end, connecting the application layers of communicating peers.
 2904 Owing to the tree topology, most connections in a Subnetwork will involve the Base Node at one end and a
 2905 Service Node at the other. However, there may be cases when two Service Nodes within a Subnetwork need
 2906 to establish connections. Such connections are called direct connections and are described in section 4.3.6.

2907 All connection establishment messages use the CON control packet. The various control packets types and
 2908 specific fields that unambiguously identify them are given in Table 22.


2909 Each successful connection established on the Subnetwork is allocated an LCID. The Base Node shall allocate
 2910 an LCID that is unique for a given LNID.

2911 **Note.** Either of the negotiating ends may decide to reject a connection establishment request. The receipt of
 2912 a connection rejection does not amount to any restrictions on making future connection requests; it may
 2913 however be advisable.



2914
 2915

Figure 90 - Connection establishment initiated by a Service Node



2916
 2917

Figure 91 - Connection establishment rejected by the Base Node

2918
2919

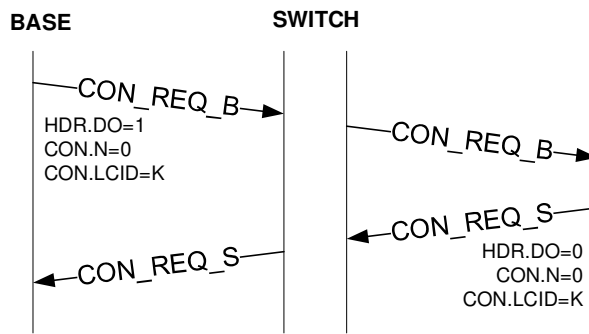


Figure 92 - Connection establishment initiated by the Base Node

2920
2921




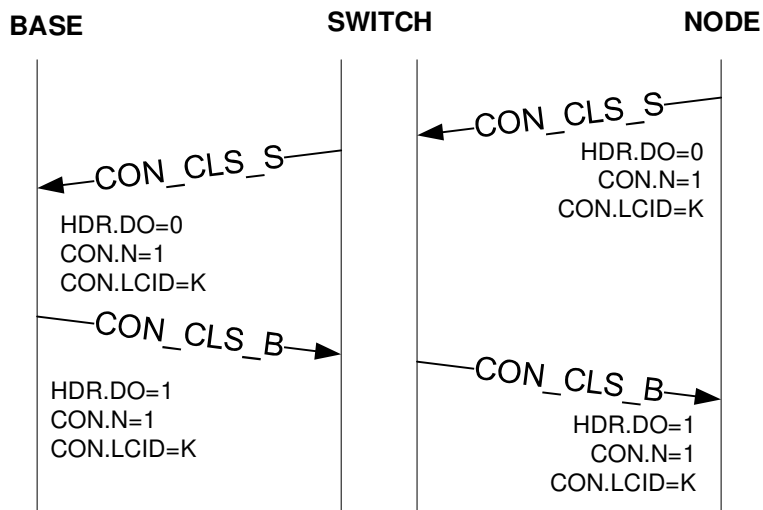
Figure 93 - Connection establishment rejected by a Service Node

2922 4.6.6.1 Connection closing

2923 Either peer at both ends of a connection may decide to close the connection at anytime. The CON control
2924 packet is used for all messages exchanged in the process of closing a connection. The relevant CON control
2925 packet fields in closing an active connection are CON.N, CON.LCID and CON.TYPE. All other fields shall be set
2926 to 0x0.

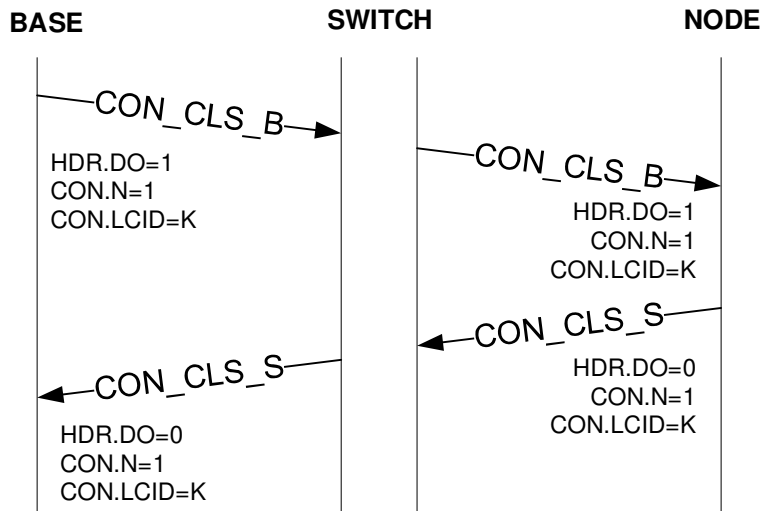
2927 A connection closure request from one end is acknowledged by the other end before the connection is
2928 considered closed. The present version of this specification does not have any explicit message for rejecting
2929 a connection termination requested by a peer at the other end.

2930 Figure 94 and Figure 95 show message exchange sequences in a connection closing process.



2931
2932

Figure 94 - Disconnection initiated by a Service Node



2933
2934

Figure 95 - Disconnection initiated by the Base Node

2935 **4.6.7 Multicast group management**

2936 **4.6.7.1 General**

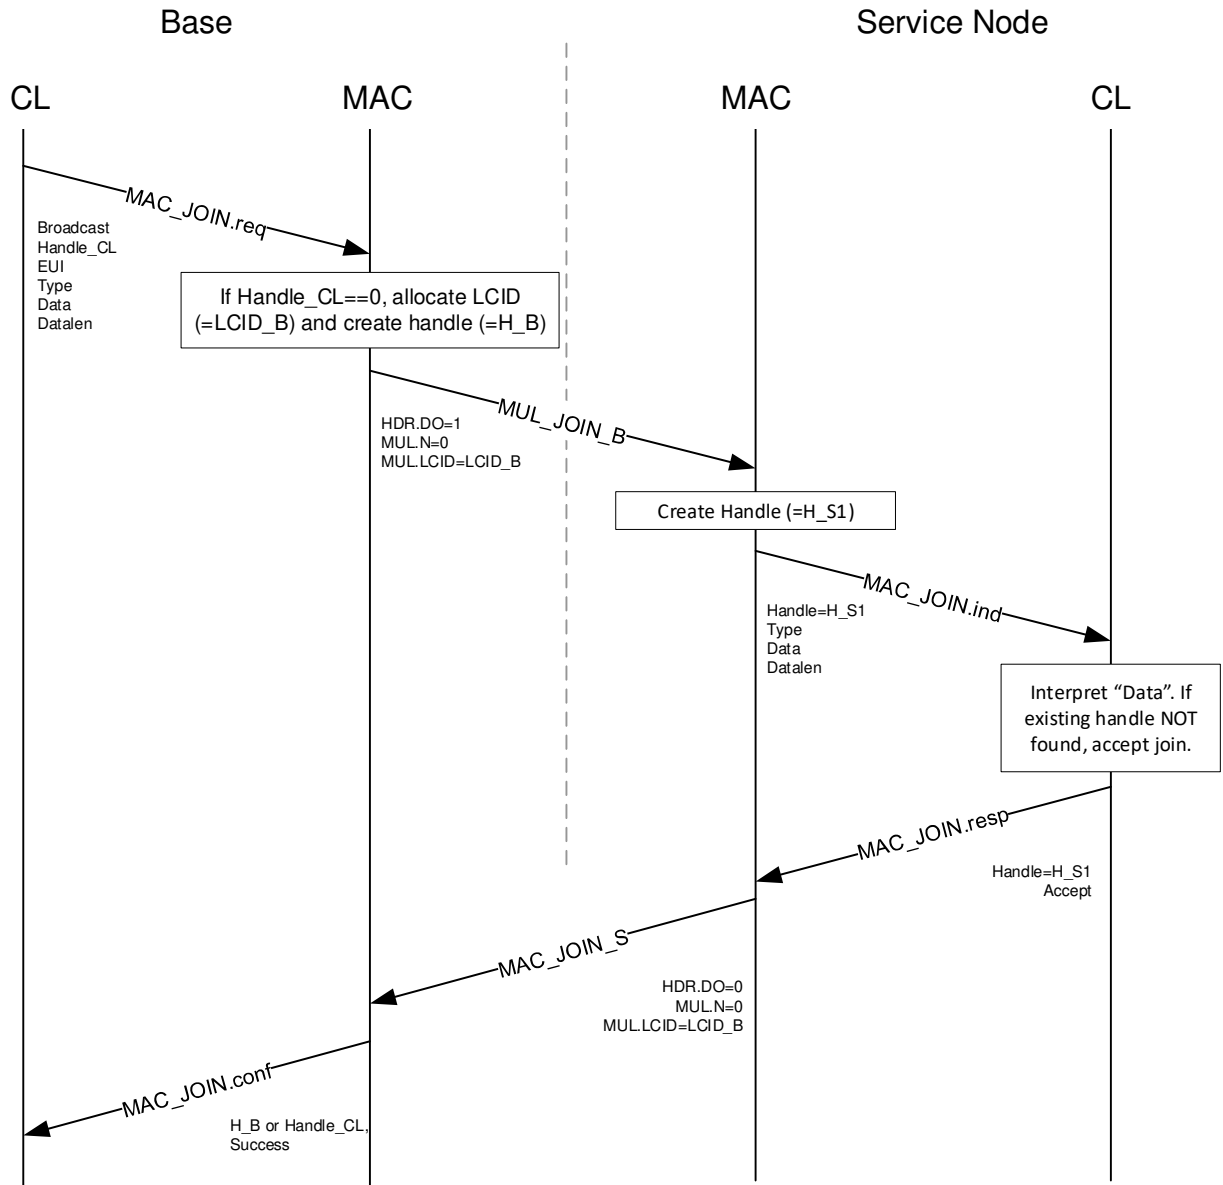
2937 The joining and leaving of a multicast group can be initiated by the Base Node or the Service Node. The MUL
2938 control packet is used for all messages associated with multicast and the usual retransmit mechanism for
2939 control packets is used. These control messages are unicast between the Base Node and the Service Node.

2940 **4.6.7.2 Group Join**

2941 Multicast group join maybe initiated from either the Base Node or Service Node. A device shall not start a
2942 new join procedure before an existing join procedure started by itself is completed.

2943 Certain applications may require the Base Node to selectively invite certain Service Nodes to join a specific
2944 multicast group. In such cases, the Base Node starts a new group and invites Service Nodes as required by
2945 application.

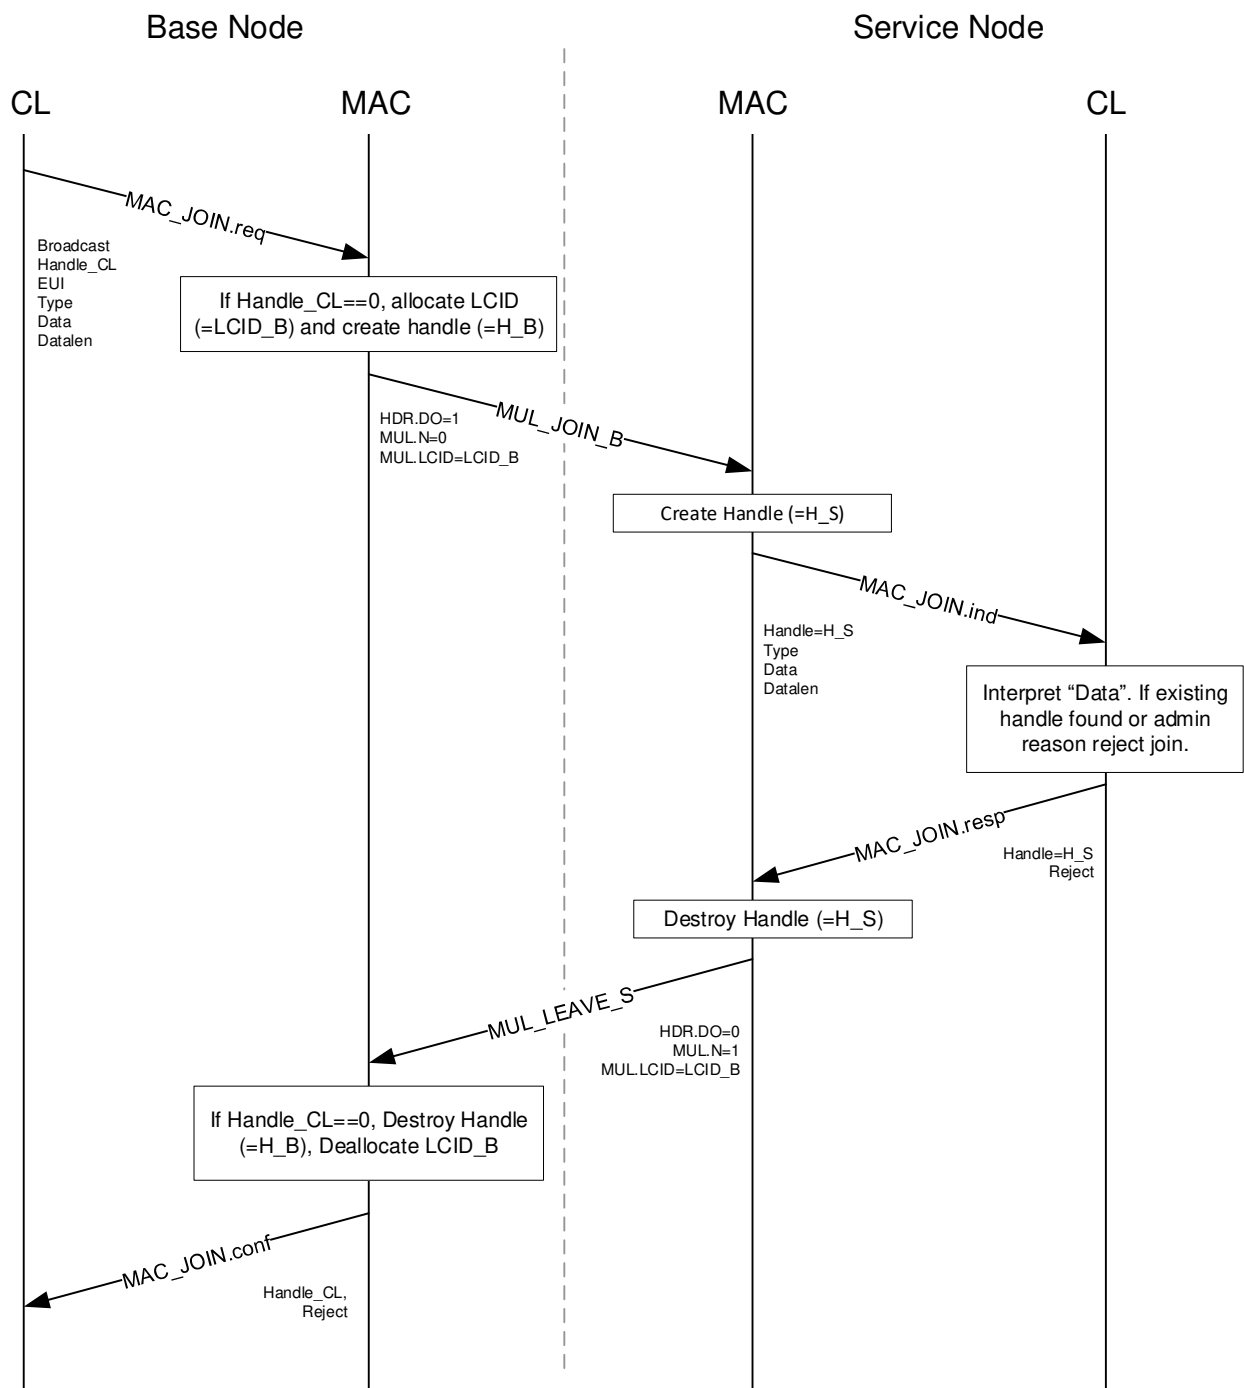
2946 Successful and failed group joins initiated from Base Node are shown in Figure 96 and Figure 97.



2947

2948

Figure 96 - Successful group join initiated by Base Node

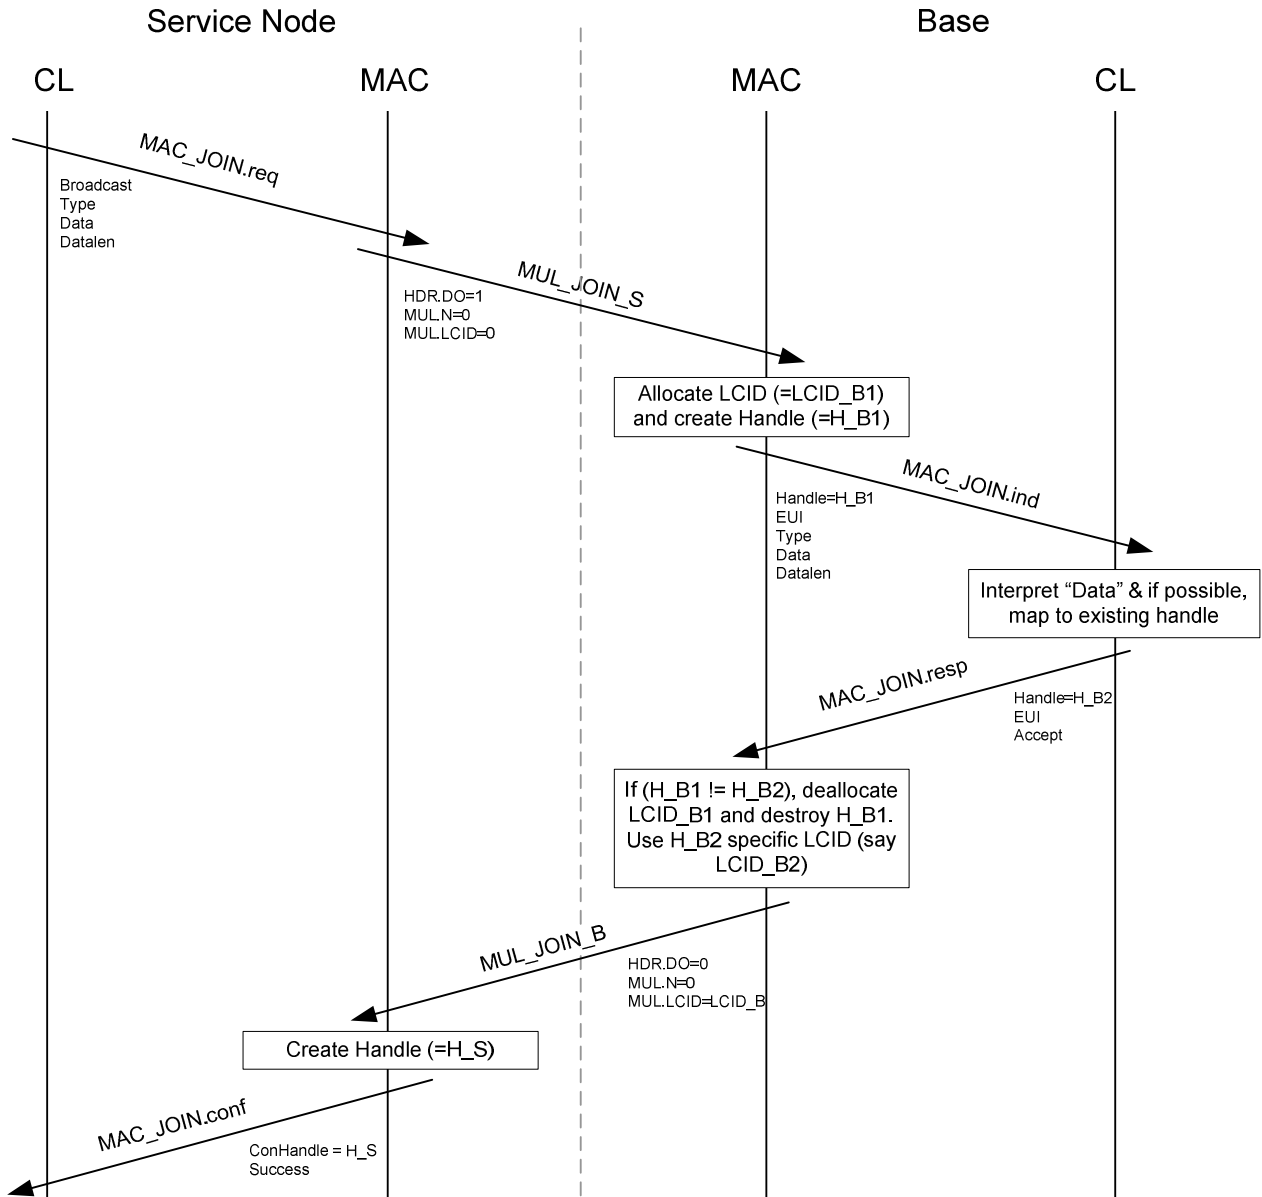


2949

2950

Figure 97 - Failed group join initiated by Base Node

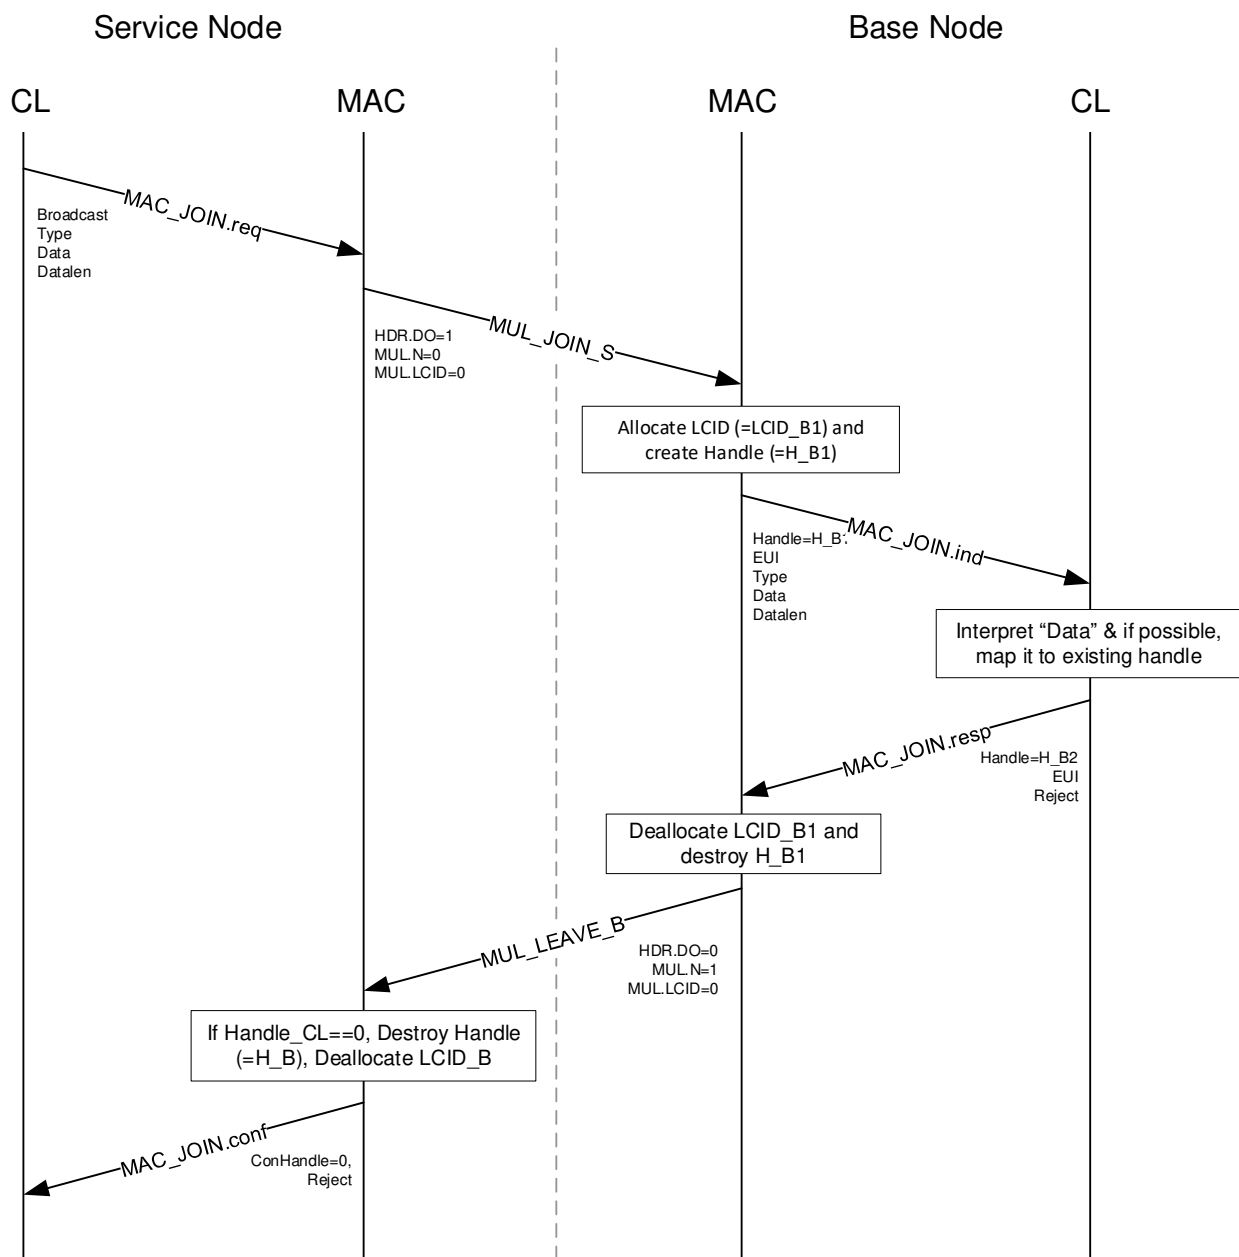
2951 Successful and failed group joins initiated from Service Node are shown in Figure 98 and Figure 99



2953

2954

Figure 98 - Successful group join initiated by Service Node



2955

2956

Figure 99 - Failed group join initiated by Service Node.

2957

4.6.7.3 Group Leave

2958

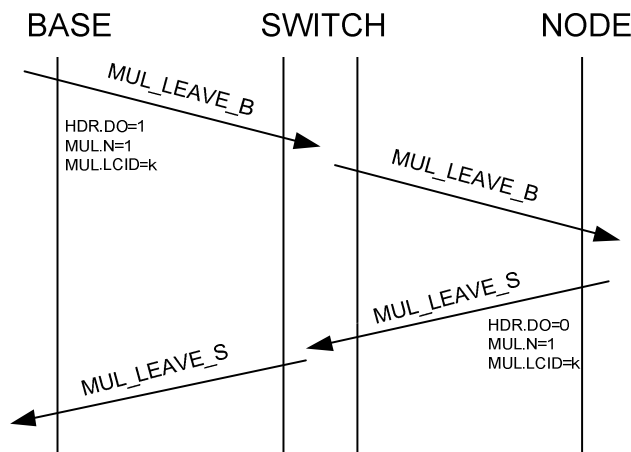
Leaving a multicast group operates in the same way as connection removal. Either the Base Node or Service

2959

Node may decide to leave the group. A notable difference in the group leave process as compared to a group

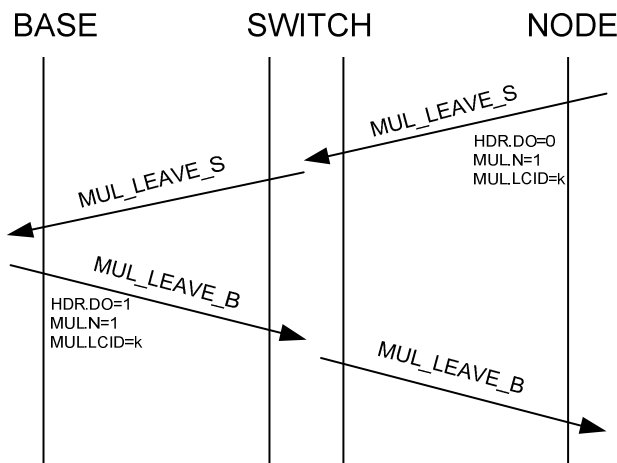
2960

join is that there is no message sequence for rejecting a group leave request.



2961
2962

Figure 100 - Leave initiated by the Base Node



2963
2964

Figure 101 - Leave initiated by the Service Node

2965 **4.6.7.4 Multicast Switching Tracking**

2966 Switch Nodes need to be aware of the multicast groups under their switching domain. Instead of having to
2967 store all the tracking information on the switches themselves, they just need to have a simple multicast table
2968 which is managed by both the Switch Node and the Base Node.

2969 Switch Nodes should just monitor muticast join operations through MUL_JOIN messages in order to start
2970 switching multicast traffic, and monitor MUL_SW_LEAVE messages in order to stop switching multicast
2971 traffic.

2972 **4.6.7.4.1 Multicast Switching Tracking for Group Join**

2973 The following rules apply for switching of traffic on a multicast group join:

- 2974 • On a successful group join from a Service Node in its control hierarchy, a Switch Node adds a new
2975 multicast Switch entry for the group LCID, where necessary. For this purpose, MUL_JOIN messages
2976 are used.
- 2977 • From that moment on, the Switch Node will switch multicast traffic for that LCID, and stops keeping
2978 track of any control message related to that group.

- The Base Node shall track all the Switch Nodes that switch multicast traffic for every multicast group. Figure 102 exemplifies the process and interactions, for the both cases of the group join initiated by the Base Node and by the Service Node and complements the processes illustrated in the figures of section 4.6.7.2.

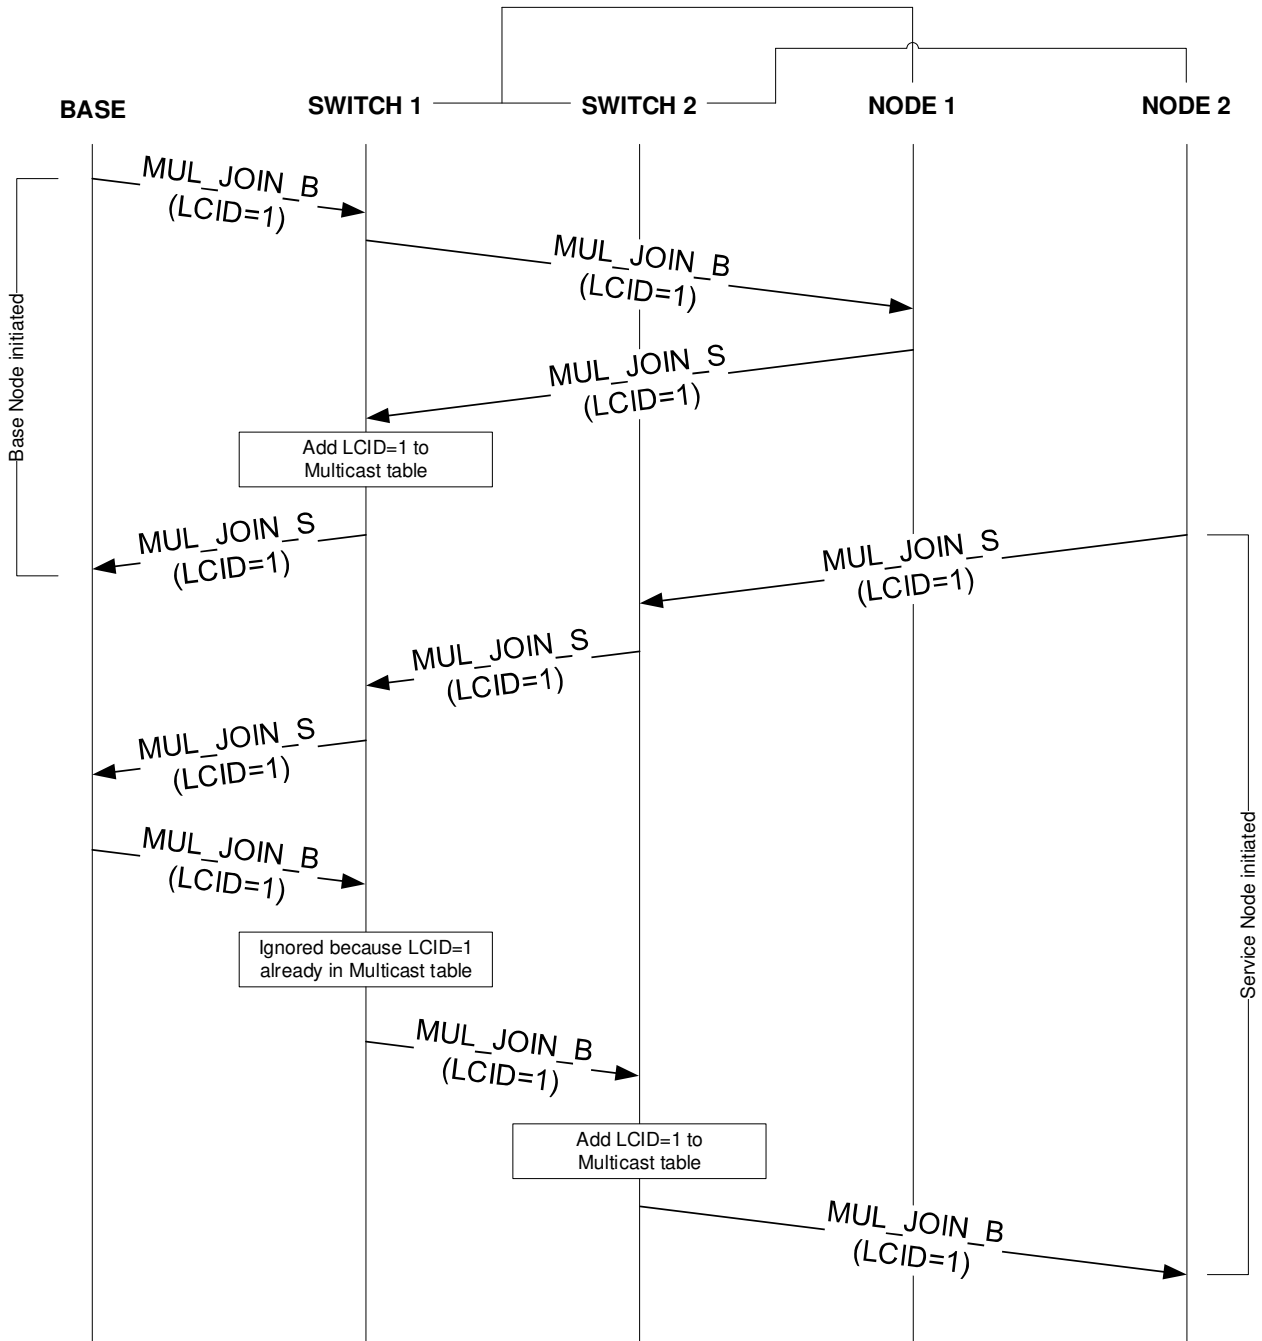


Figure 102 - Multicast Switching Tracking for Group Join.

2982

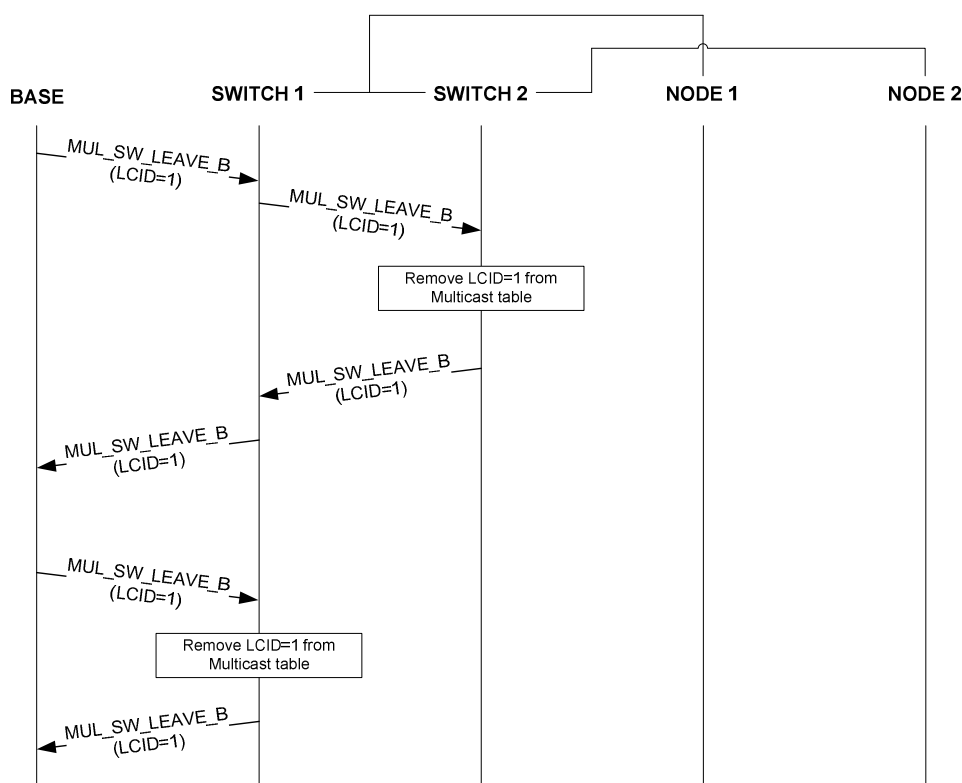
2983

4.6.7.4.2 Multicast Switching Tracking for Group Leave

For switching of traffic on a multicast group leave, the Base Node shall monitor when all nodes depending on a Switch Node leave a given multicast group, and start a `MUL_SW_LEAVE` procedure to remove that multicast entry for that Switch Node and group.

2987

2988 Figure 103 exemplifies the process and interactions; when they are executed, they take place after the
 2989 processes illustrated in the figures of section 4.6.7.3.



2990
 2991 **Figure 103 - Multicast Switching Tracking for Group Leave.**

2992 **4.6.7.4.3 Multicast Switching with double switching**

2993 A Switch Node that is transmitting BPDUs on both Type A and Type B PHY frames can switch all data arriving
 2994 on a multicast connection using both types of PHY frames. This implies replicating data while transmitting
 2995 twice but this will enable coverage across its entire control domain. Future versions of this specification can
 2996 further optimize on this to avoid some unwanted traffic that maybe generated by taking this generic
 2997 approach. This version leaves it open for implementations to either optimize their decision process or replicate
 2998 data using both modulation schemes.

2999 No matter the policy implemented by the double switch, the device that receives the same broadcast or
 3000 multicast packet with both modulations shall be able to discard one replica (see section 4.6.3.2).

3001 On receipt of a MUL_SW_LEAVE message, the Switch Node shall stop further switching of multicast data for
 3002 the corresponding connection, on both PHY frame types.

3003 **4.6.8 Robustness Management**

3004 **4.6.8.1 General**

3005 The Robustness-management (RM) mechanism is designed to select the most suitable transmission scheme
 3006 from the eight available ones (Robust DBPSK, Robust DQPSK, DBPSK_CC, DBPSK, DQPSK_CC, DQPSK,

3007 D8PSK_CC and D8PSK). Depending on the transmission channel conditions, the nodes shall decide either to
3008 increase the robustness or to select faster transmission modes.

3009 Note that the mechanism described here shall be used to decrease and increase the robustness of Generic
3010 DATA packets. MAC control packets shall be transmitted in conformance with specification in Section
3011 4.3.3.3.3.

3012 By default, decision about applicable transmission mode is taken locally. That is, dynamic adaptation of the
3013 transmission mode is performed taking into account link level channel information, which is exchanged
3014 between any pair of nodes in direct vision (parent and child). As an exception to this rule, a Base Node may
3015 decide to disable dynamic robustness-management and force a specific transmission mode in the Service
3016 Node(s). This static configuration shall be fixed during registration, as explained in 4.6.1.

3017 The robustness-management mechanism comprises two main features:

- 3018 • Link quality information embedded in the packet header of any Generic packets
- 3019 • Link level ACK-ed ALIVE mechanism, as explained in 4.6.5.

3020 4.6.8.2 Link quality information embedded in the packet header

3021 All Generic packets shall convey link quality related information. Four bits in the packet header - “PKT.RM”,
3022 see 4.4.2.3 – are used by the transmitting device to notify the other peer of the weakest modulation scheme
3023 that the transmitter considers it could receive. The transmitting device calculates this value processing the
3024 received packets sent by the other peer. The calculation of PKT.RM value is implementation dependent.

3025 Whenever a node receives a Generic packet from a peer, it shall update the peer related info contained in
3026 *macListPhyComm* PIB as follows:

- 3027 • Store “PKT.RM” from the received packet in *macListPhyComm.phyCommTxModulation*
- 3028 • Reset *macListPhyComm.phyCommRxAge* (time [seconds] since the last update of
3029 *phyCommTxModulation*).

3030 Whenever a node wants to transmit DATA to an existing peer, it shall check validity of the robustness-
3031 management information it stores related to that peer. The maximum amount of time that robustness-
3032 management information is considered to be valid without any further update is specified in PIB
3033 *macUpdatedRMTimeout*. Consequently, the node shall compare *phyCommRxAge* (time since the last update)
3034 with *macUpdatedRMTimeout*:

- 3035 • *macListPhyComm.phyCommRxAge* \geq *macUpdatedRMTimeout*: The node shall transmit using the
3036 most robust modulation scheme available for the PHY frame type in use. **Note:** the first time a node
3037 sends DATA to one peer, RM information is automatically considered to be “out of date” and
3038 consequently the most robust modulation scheme available shall be used.
- 3039 • *macListPhyComm.phyCommRxAge* $<$ *macUpdatedRMTimeout*: The node shall check the value
3040 stored in the *phyCommTxModulation* field:
 - 3041 ○ Different from 0xF: The modulation to be used shall be the same as specified in
3042 *phyCommTxModulation*.
 - 3043 ○ Equal to 0xF: The node shall transmit using the most robust modulation scheme available
3044 for the PHY frame type in use.

3045 4.6.8.3 Link level ACK-ed ALIVE mechanism

3046 Alive procedure defines repetitions that are performed in every hop as described in 4.6.5. An
3047 ALV_REQ_B/ALV_RSP_S transmitting device shall use this fact to assume a delivery failure if it does not
3048 receive the corresponding ACK packet. In this case the transmitting device shall re-transmit the packet: the
3049 first repetition shall be performed with the same robustness, which will be successively increased after
3050 every link level repetition. Once the maximum number of repetitions is reached, the most robust
3051 modulation in which the node can transmit shall be stored for that link, even if the repetitions were due to
3052 the ACK packets.

3053 The device receiving the ALV_REQ_B/ALV_RSP_S, on reception of a packet being sent more than twice
3054 ($ALV.TX_SEQ > 1$), shall send the ACK packet with at least the same robustness as the received packet.

3055 The ALV packets shall be transmitted in one of the following encodings: DBPSK_CC, Robust DQPSK and
3056 Robust DBPSK. The robustness increase should be performed in that order.

3057 In the Terminal Node a three way handshake is performed, once the ALV_REQ_B has arrived the Service
3058 Node shall start a regular ALV_RSP_S send transaction following the same rules.

3059 In every case the ALV.RX_SNR, ALV.RX_POW and ALV.RX_ENC shall send those PHY parameters of the last
3060 received ALV_REQ_B/ALV_RSP_S packet, and in the PKT.RM they shall send the least robust modulation in
3061 which it should be able to receive.

3062 4.6.8.4 PHY robustness changing

3063 From the PHY point of view there are several parameters that may be adjusted and which affect the
3064 transmission robustness: the transmission power and modulation parameters (convolutional encoding and
3065 constellation). As a general rule the following rules should be followed:

- 3066 • **Increase robustness:** increase the power and, if it is not possible, improve the modulation scheme
3067 robustness (reducing throughput).
- 3068 • **Reduce robustness:** reduce the modulation scheme robustness (increasing throughput) and, if it is
3069 not possible, reduce the transmission power.

3070 4.6.9 Channel allocation

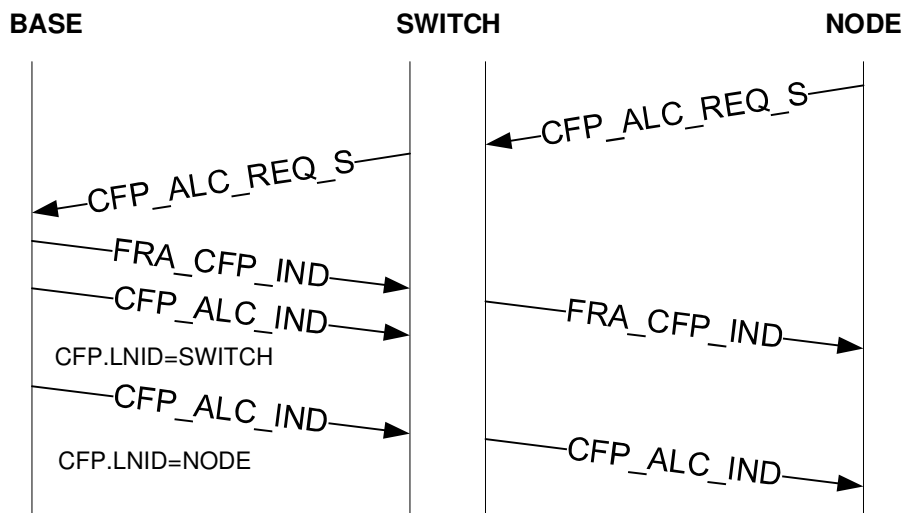
3071 Allocation of specific channel resources is possible in the CFP. Each MAC frame shall include a contention free
3072 period with a minimum duration of $(MACBeaconLength1 + 2 \times macGuardTime)$, which may be used for
3073 beacon transmission and/or allocation of specific application data transmissions. Any kind of CFP usage,
3074 either beacon transmission or allocation of channel resources for data, shall be always granted by the Base
3075 Node.

3076 4.6.9.1 Beacon channel allocation

3077 As part of a promotion procedure, a Terminal node may be promoted to Switch status and gain the ability to
3078 transmit its own beacons. These beacons shall be allocated in the CFP as explained in 4.6.3.

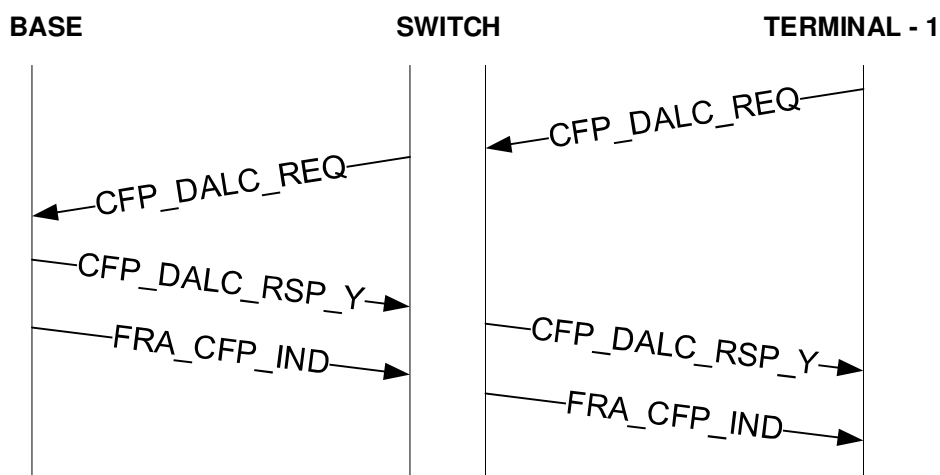
3079 **4.6.9.2 Data channel allocation**

3080 A CFP allocation / de-allocation request to transport application data may be initiated either by the Base
 3081 Node or the Service Node. The CFP MAC control packet described in 4.4.2.6.7 shall be used for that purpose.
 3082 Figure 104 below shows a successful channel allocation sequence. All channel allocation requests initiated
 3083 by Service Node are forwarded to the Base Node. Note that in order to assure a contention-free channel
 3084 allocation along the entire path, the Base Node allocates non-overlapping times to intermediate Switch
 3085 Nodes. In a multi-level Subnetwork, the Base Node may also reuse the allocated time at different levels.
 3086 While reusing the said time, the Base Node needs to ensure that the levels that use the same time slots have
 3087 sufficient separation so that there is no possible interference.



3088
 3089 **Figure 104 - Successful allocation of CFP period**

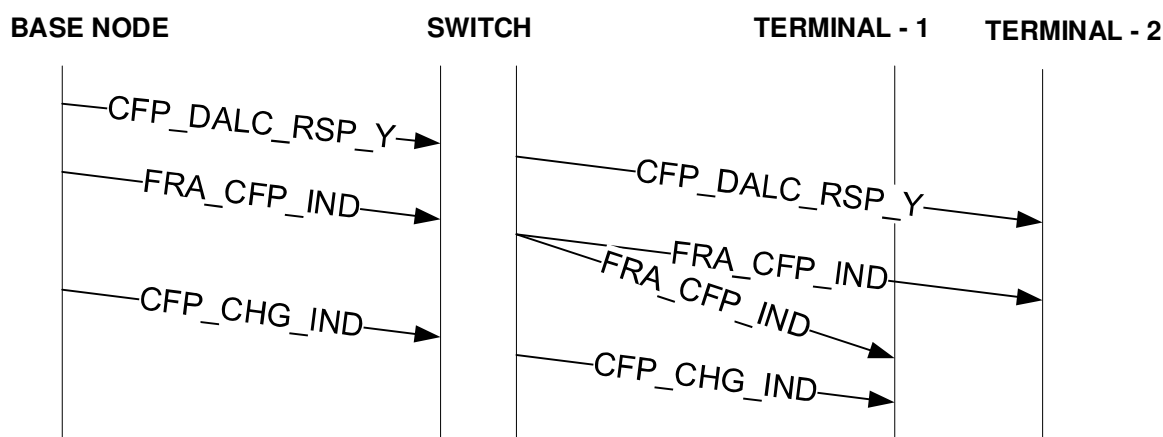
3090 Figure 105 below shows a channel de-allocation request from a Terminal device and the resulting
 3091 confirmation from the Base Node.



3092
 3093 **Figure 105 - Successful channel de-allocation sequence**

3094 Figure 106 below shows a sequence of events that may lead to a Base Node re-allocation contention-free
 3095 slot to a Terminal device that already has slots allocated to it. In this example, a de-allocation request from
 3096 Terminal-2 resulted in two changes: firstly, in global frame structure, this change is conveyed to the

3097 Subnetwork in the FRA_CFP_IND (a standard FRA packet intended to change CFP duration only) packet;
3098 secondly, it is specific to the time slot allocated to Terminal-1 within the CFP.



3099
3100

Figure 106 - Deallocation of channel to one device results in the change of CFP allocated to another

3101 4.7 Automatic Repeat Request (ARQ)

3102 4.7.1 General

3103 Devices complying with this specification may either implement an ARQ scheme as described in this section
3104 or no ARQ at all. This specification provides for low-cost Switch and Terminal devices that choose not to
3105 implement any ARQ mechanism at all.

3106 4.7.2 Initial negotiation

3107 ARQ is a connection property. During the initial connection negotiation, the originating device indicates its
3108 preference for ARQ or non-ARQ in CON.ARQ field. The responding device at the other end can indicate its
3109 acceptance or rejection of the ARQ in its response. If both devices agree to use ARQ for the connection, all
3110 traffic in the connection will use ARQ for acknowledgements, as described in Section 4.7.3. If the responding
3111 device rejects the ARQ in its response, the data flowing through this connection will not use ARQ.

3112 4.7.3 ARQ mechanism


3113 4.7.3.1 General

3114 The ARQ mechanism works between directly connected peers (original source and final destination), as long
3115 as both of them support ARQ implementation. This implies that even for a connection between the Base
3116 Node and a Terminal (connected via one or more intermediate Switch devices), ARQ works on an end-to-end
3117 basis. The behavior of Switch Nodes in an ARQ-enabled connection is described in Section 4.7.4. When using
3118 ARQ, a unique packet identifier is associated with each packet, to aid in acknowledgement. The packet
3119 identifier is 6 bits long and can therefore denote 64 distinct packets. ARQ windowing is supported, with a
3120 maximum window size of 32 (5 bits), as described in Section 4.7.3.3.

3121 **4.7.3.2 ARQ PDU**


3122 **4.7.3.2.1 General**

3123 The ARQ subheader contains a set of bytes, each byte containing different subfields. The most significant bit
 3124 of each byte, the M bit, indicates if there are more bytes in the ARQ subheader.




3125
 3126 **Figure 107 - ARQ subheader only with the packet id**

3127 Figure 107 shows an ARQ subheader with the first M bit of 0 and so the subheader is a single byte
 3128 and contains only the packet ID for the transmitted packet.




3129
 3130 **Figure 108 - ARQ subheader with ARQ.INFO**


3131 Figure 108 has the M bit in the first byte of the ARQ subheader set, and so the subheader contains multiple
 3132 bytes. The first byte contains the packet ID of the transmitted packet and then follows the ARQ.INFO which
 3133 is a list of one or more bytes, where each byte could have one of the following meanings:



3134
 3135 **Figure 109 - ARQ.ACK byte fields**



3136
 3137 **Figure 110 - ARQ.WIN byte fields**



3138
 3139 **Figure 111 - ARQ.NACK byte fields**

3140 If there are multiple packets lost, an ARQ.NACK is sent for each of them, from the first packet lost to the last
 3141 packet lost. When there are several ARQ.NACK they implicitly acknowledge the packets before the first
 3142 ARQ.NACK, and the packets in between the ARQ.NACKs. If an ARQ.ACK is present, it shall be placed at the
 3143 end of the ARQ subheader, and shall reference to an ARQ.ACKID that is later than any other ARQ.NACKID, if
 3144 present. If there is at least an ARQ.NACK and an ARQ.ACK they also implicitly acknowledge any packet in the
 3145 middle between the last ARQ.NACKID and the ARQ.ACK.

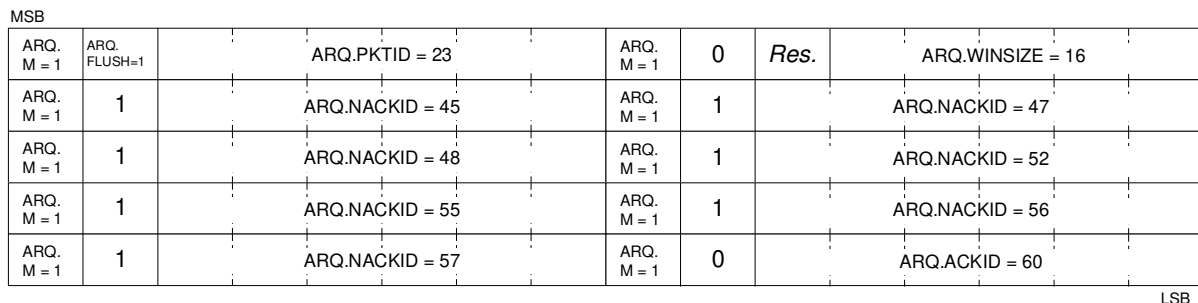
3146 For interoperability, a device shall be able to receive any well-formed ARQ subheader and shall process at
 3147 least the first ARQ.ACK or ARQ.NACK field.

3148 The subfields have the following meanings as described in Table 56:

3149 **Table 56 - 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.

3150 **4.7.3.2.2 ARQ subheader example**



3151

3152 **Figure 112 - Example of an ARQ subheader with all the fields present**

3153 In this example all the ARQ subheader fields are present. To make it understandable, since both Nodes are
 3154 both transmitters and receivers, the side receiving this header will be called A and the other side transmitting
 3155 B. The message has the packet ID of 23 if it contains data; otherwise the next data packet to be sent has the
 3156 packet ID of 23. Since the flush bit is set it needs to be ACKed/NACKed.

3157 B requests the retransmission of packets 45, 47, 48, 52, 55, 56 and 57. ACK = 60, so it has received packets
 3158 <45, 46, 49, 50, 51, 53, 54, 58 and 59.

3159 The window is 16 and it has received and processed up to packet 44 (first NACK = 45), so A can send all
 3160 packets <= 60; that is, as well as sending the requested retransmits, it can also send packet ID = 60.

3161 **4.7.3.3 Windowing**

3162 A new connection between two peer devices starts with an implicit initial receiver window size of 1 and a
 3163 packet identifier 0. This window size is a limiting case and the transaction (to start with) shall behave like a
 3164 “Stop and Wait” ARQ mechanism.

3165 On receipt of an ARQ.WIN, the sender would adapt its window size to *ARQ.WINSIZE*. This buffer size is
 3166 counted from the first packet completely ACK-ed, so if there is a NACK list and then an ACK the window size
 3167 defines the number of packets from the first NACK-ed packet that could be sent. If there is just an ACK in the
 3168 packet (without any NACK) the window size determines the number of packets that can be sent from that
 3169 ACK.

3170 An *ARQ.WINSIZE* value of 0 may be transmitted back by the receiver to indicate congestion at its end. In such
 3171 cases, the transmitting end should wait for at least *ARQCongClrTime* before re-transmitting its data.

3172 **4.7.3.4 Flow control**

3173 The transmitter must manage the ACK sending algorithm by the flush bit; it is up to it having a proper ARQ
 3174 communication. The receiver is only forced to send ACKs when the transmitter has sent a packet with the
 3175 flush bit set, although the receiver could send more ACKs even if not forced to do it, because the flow control
 3176 is only a responsibility of the transmitter. The transmitter shall close the connection latest if the Packet
 3177 acknowledgement is missing after *ARQMaxTxCount* Packet retransmissions. The transmitter may choose to
 3178 use lower maximum retransmit value than *ARQMaxTxCount* and it may also close the connection any time
 3179 earlier if it determines proper data exchange cannot be restored.

3180 These are the requisites to be interoperable, but the algorithm is up to the manufacturer. It is strongly
 3181 recommended to piggyback data-ACK information in outgoing packets, to avoid the transmission of
 3182 unnecessary packets just for ACK-ing. In particular in order to allow consolidated ACKs or piggybacking, the
 3183 maximum time for each implementation before sending an ACK is *ARQMaxAckHoldTime*.

3184 **4.7.3.5 Algorithm recommendation**


3185 No normative algorithm is specified, for a recommendation see Annex I.

3186 **4.7.3.6 Usage of ARQ in resource limited devices**

3187 Resource limited devices may have a low memory and simple implementation of ARQ. They may want to use
 3188 a window of 1 packet. They work as a “Stop and Wait” mechanism.


3189 The ARQ subheader to be generated shall be one of the followings:

3190 If there is nothing to acknowledge:




3191
 3192 **Figure 113 - Stop and wait ARQ subheader with only packet ID**

3193 If there is something to acknowledge carrying data:



3194
3195 **Figure 114 - Stop and wait ARQ subheader with an ACK**

3196 If there is something to acknowledge but without any data in the packet:



3197
3198 **Figure 115 - Stop and wait ARQ subheader without data and with an ACK**

3199 The ARQ.WINSIZE is not generally transmitted because the window size is already 1 by default, it only may
3200 be transmitted to handle congestion and to resume the transmission again.

3201 **4.7.4 ARQ packets switching**

3202 All Switch Nodes shall support transparent bridging of ARQ traffic, whether or not they support ARQ for their
3203 own transmission and reception. In this mode, Switch Nodes are not required to buffer the packets of the
3204 ARQ connections for retransmission.

3205 Some Switch Nodes may buffer the packets of the ARQ connections, and perform retransmission in
3206 response to NACKs for these packets if the subnetwork is working with security profile 0. If the subnetwork
3207 is working with security profile 1 or 2, the switch node shall not perform retransmission in response to NACKs
3208 for these packets. The following general principles shall be followed.

- 3209 • The acknowledged packet identifiers shall have end-to-end coherency.
- 3210 • The buffering of packets in Switch Nodes and their retransmissions shall be transparent to the source
3211 and Destination Nodes, i.e., a Source or Destination Node shall not be required to know whether or
3212 not an intermediate Switch has buffered packets for switched data.

3213 **4.8 Time Reference**

3214 Packets in PRIME may interchange time references by providing a TREF subheader. Due to the frame and
3215 superframe structure of the MAC layer, when a node is registered to a PRIME subnetwork it is already
3216 synchronized. The TREF subheader includes a time reference that is relative to the beginning of a frame in
3217 order to make reference to a specific moment in time.

3218

Table 57 - Time Reference subheader fields

Name	Length	Description
TREF.SEQ	5 bits	Sequence number of the MAC Frame that is used as reference time of the event to notify.
Reserved	3 bits	Always 0 for this version of the specification. Reserved for future use.
TREF.TIME	32 bits	Signed number in 10s of microseconds between the moment of the event, and the beginning of the frame. Positive for events after the beginning of the MAC frame, and negative for events before the beginning of the MAC frame. 0x80000000 is a special value that means that it is an invalid time reference.

3219 During the transmission of a new packet, the transmitter of a TREF subheader should always keep the
3220 TREF.SEQ as updated as possible.

3221 During the switching of a packet with a TREF subheader, the Switch Node may update the TREF.TIME and
3222 TREF.SEQ to change the MAC frame this reference is based on, as long as it makes reference to the same
3223 instant in time. A switch shall update the fields if $(RX_SEQ - TREF.SEQ) \& 31 > (TX_SEQ - TREF.SEQ) \& 31$,
3224 where RX_SEQ is the frame sequence number when the packet is received by the switch and TX_SEQ is the
3225 sequence number when a packet is transmitted by the switch. In this case, this field may be easily updated
3226 by subtracting the length of a superframe to the TREF.TIME field, leaving the same TREF.SEQ. This mechanism
3227 is in order to avoid a superframe overlapping.

3228 If the time reference cannot be represented in the TREF.TIME field, then the field TREF.TIME should have the
3229 value 0x80000000 that means that it is an invalid reference.

3230 4.9 Backward Compatibility with PRIME 1.3.6


3231 In order to interoperate with Service Nodes conforming to v1.3.6 of specifications, v1.4 conformant devices
3232 can implement a backward-compatibility mode. Since PRIME v1.3.6 Service Nodes will not understand v1.4
3233 message formats, v1.4 compliant devices implementing backward-compatibility mode shall support
3234 additional messaging capabilities that enable them to communicate with PRIME v1.3.6 devices. Any
3235 Subnetwork that allows registration of one or more PRIME v1.3.6 device/s shall be termed to be running in
3236 "backward-compatibility" mode and will operate with the following characteristics:

- 3237 • Base Node shall always be a v1.4 compliant implementation i.e. a PRIME v1.3.6 Base Node is
3238 incapable of managing a Subnetwork in backward-compatibility mode.
- 3239 • All robust mode PDUs shall be transmitted using PHY BC Frames as defined in Annex K.
- 3240
- 3241 • To accommodate for size restrictions in PHY BC Frames, the Base Node shall limit the maximum SAR
3242 segment size to be less or equal than 64 bytes for all service nodes located, directly or indirectly,
3243 behind a robust link

3244 **4.9.1 Frame Structure and Channel Access**

3245 A Subnetwork in “backward-compatibility” mode shall abide by principles laid down in points below:


- 3246 • Fixed frame length of 276 symbols shall be used. The frames include up to five consecutive non-robust beacon slots, each of them having a length of 11.008ms.
- 3247
- 3248 • Transition to longer frames is prohibited. The base node transmits a beacon using DBPSK_CC modulation in every frame in beacon slot 0. The frame format is shown in Figure 116.
- 3249



3250

3251 **Figure 116 - CBCN Frame format for backwards compatibility mode**

- 3252 • PRIME v1.4 compatibility mode allows the Base Node to place up to two robust mode beacons per frame at the end of the CBCN Frame. They are located at the end of CFP, respecting the guard times (macGuardTime) before each one and at the end of the frame, as shown in the Figure 117.
- 3253
- 3254 • Robust beacon allocation policy is up to the manufacturer, it could be 0 to 2 robust beacons and can be allocated from the beginning or on demand.
- 3255
- 3256 • The robust beacons shall be sent in DBPSK_R encoding with the PHY frame type BC.
- 3257



3258

3259 **Figure 117 - Robust beacon allocation in 1.4 BC**

- 3260 • The Base Node shall set the CFP duration to a value which guarantees that the robust mode beacons are fully located within CBCN.CFP.
- 3261
- 3262 • For a Switch Node using DBPSK_CC for beacon transmission, the Base Node allocates space from the non-robust beacon slots (slot 0 to 4). Beacon slot allocation rules according to PRIME v1.3.6 shall apply for allocation of non-robust beacon slots.
- 3263
- 3264 • For allocation of robust beacons the Base Node should not update the beacon slot count, as the robust beacons shall be placed in the CFP.
- 3265
- 3266 • The Base Node can also transmit robust beacons but it is not mandated to do so.
- 3267
- 3268 • If a switch or Base Node transmits robust beacons, it shall transmit at least once every 32 sub-frames.
- 3269 • Frame format shall respect the same restrictions on SCP and CFP durations defined in PRIME v1.3.6.
- 3270 • CSMA/CA algorithm defined in 1.3.6 shall be used.

3271 **4.9.2 Switching**

3272 In a network running in PRIME v1.4 compatibility mode, uplink (HDR.DO=0) Multicast and Broadcast packets
3273 shall follow the same rules as the ones described in PRIME version 1.3.6:

- 3274 • The node shall only switch the packet that has a broadcast or multicast destination (PKT.LNID = 0x3FFF
3275 or 0x3FFE) and that was transmitted by a Node registered through this Switch Node (PKT.SID=LSID of
3276 this Switch Node).
- 3277 • In case a broadcast or multicast packet is switched up, the Switch Node shall replace the PKT.SID with
3278 Switch Node's SID.

3279 For the rest of the packet types the Switch Node shall follow the rules and actions described in chapter 4.3.5.

3280 **4.9.3 PDU Frame Formats**

3281 **4.9.3.1 General Format**

3282 In a network running in PRIME v1.4 compatibility mode, all nodes shall use the standard Generic Mac
3283 Header (see Section 4.4.2.4) and the Compatibility Packet Header (CPKT, see Annex K). The CRC calculation
3284 follows the standard procedure described in Section 4.4.2. These headers and CRC calculation follow the
3285 PRIME v1.3.6 specification.

3286 In a compatibility mode network, some control messages need a different format from the standard PRIME
3287 v1.4 format. This is for example the case for messages which are sniffed by PRIME v1.3.6 devices. These
3288 special messages are listed in the following sub-sections. On the other hand, the standard PRIME v1.4
3289 payload format is used for the CON, CFP, MUL and SEC control packets.

3290 **4.9.3.2 Registration and Unregistration control messages**

3291 A mixture of compatibility mode registration messages (CREG, Annex K.1.1.1.1), which follow the following
3292 PRIME v1.3.6 message format, and PRIME v1.4 format REG control messages shall be used. For a detailed
3293 description of the registration procedure in a compatibility mode network see Annex K.2.1.

3294 The messages used during unregistration shall follow the CREG frame format specified in Annex K.1.1.1.1.

3295 **4.9.3.3 Promotion and Beacon Slot Indication control messages**

3296 For all promotion messages the compatibility mode format CPRO (see Annex K.1.1.1.2) shall be used. The
3297 CREG messages resemble PRIME v1.3.6 messages. This is important as switches on the branch need to sniff
3298 these packets in order to refresh their switch tables. The compatibility mode promotion procedure, which is
3299 described in Annex K.2.3, requires also compatibility BSI packets (CBSI). The BSI packets are no longer used
3300 in PRIME v1.4. A compatibility mode network does not support a modulation change of a switch.

3301 The messages used during demotion shall follow the CPRO frame format specified in Annex K.1.1.1.2.

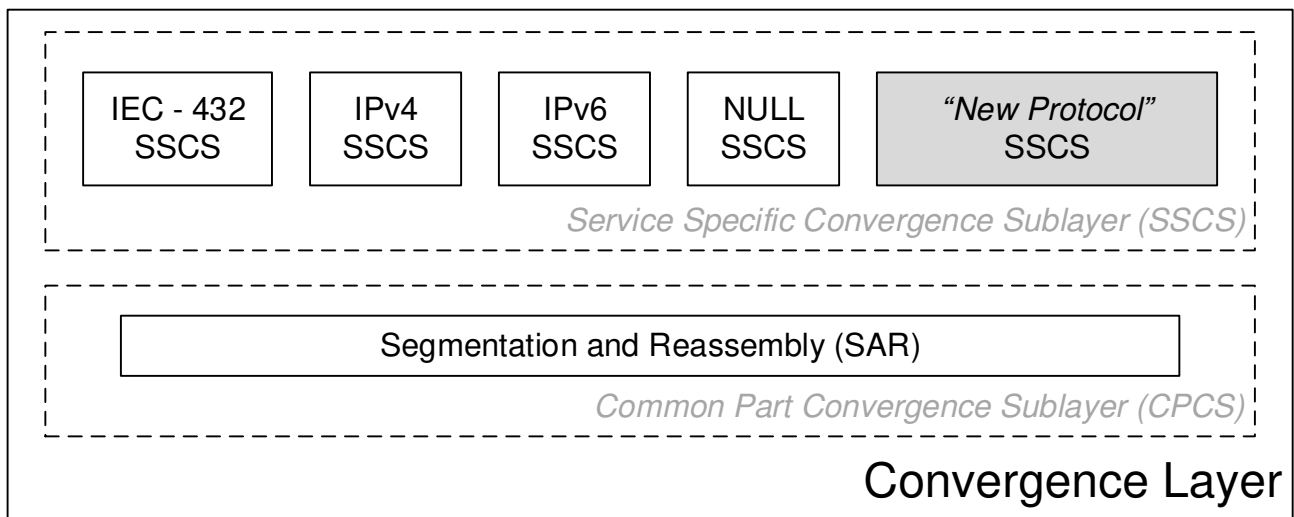
3302 **4.9.3.4 Keep-Alive control messages**

3303 PRIME v1.3.6 service nodes do not know about the new link level keep-alive process. Therefore, for networks
3304 operating in backward compatibility mode, the keep-alive process shall remain the same as in PRIME v1.3.6.
3305 The process is described in Annex K.2.5. In addition, the CALV control messages are also enumerated in
3306 K.1.1.1.5.

3307 5 Convergence layer

3308 5.1 Overview

3309 Figure 118 shows the overall structure of the Convergence layer.



3310

3311

Figure 118 - Structure of the Convergence layer

3312 The Convergence layer is separated into two sublayers. The Common Part Convergence Sublayer (CPCS)
 3313 provides a set of generic services. The Service Specific Convergence Sublayer (SSCS) contains services that
 3314 are specific to one communication profile. There are several SSCSs, typically one per communication profile,
 3315 but only one CPCS. The use of CPCS services is optional in that a certain SSCS will use the services it needs
 3316 from the CPCS, and omit services which are not needed.

3317 5.2 Common Part Convergence Sublayer (CPCS)

3318 5.2.1 General

3319 This specification defines only one CPCS service: Segmentation and Reassembly (SAR).

3320 5.2.2 Segmentation and Reassembly (SAR)


3321 5.2.2.1 General

3322 CPCS SDUs which are larger than 'macSARSize-1' bytes are segmented at the CPCS. CPCS SDUs which are
 3323 equal or smaller than 'macSARSize -1' bytes may also optionally be segmented. Segmentation means
 3324 breaking up a CPCS SDU into smaller parts to be transferred by the MAC layer. At the peer CPCS, the smaller
 3325 parts (segments) are put back together (i.e. reassembled) to form the complete CPCS SDU. All segments
 3326 except the last segment of a segmented SDU must be the same size and at most macSARSize bytes in length.

3327 Segments may be decided to be smaller than ‘macSARSize -1’ bytes e.g. when the channel is poor. The last
 3328 segment may of course be smaller than ‘macSARSize -1’ bytes.

3329 In order to keep SAR functionality simple, the macSARSize is a constant value for all possible
 3330 modulation/coding combinations at PHY layer. The value of macSARSize is such that with any
 3331 modulation/coding combination, it is always possible to transmit a single segment in one PPDU. Therefore,
 3332 there is no need for discovering a specific MTU between peer CPCs or modifying the SAR configuration for
 3333 every change in the modulation/coding combination. In order to increase efficiency, a Service Node which
 3334 supports packet aggregation may combine multiple segments into one PPDU when communicating with its
 3335 peer.

3336 Segmentation always adds a 1-byte header to each segment. The first 2 bits of SAR header identify the type
 3337 of segment. The semantics of the rest of the header information then depend on the type of segment. The
 3338 structure of different header types is shown in Figure 119 and individual fields are explained in Table 52. Not
 3339 all fields are present in each SAR header. Either SAR.NSEGS or SAR.SEQ is present, but not both.



3340
 3341 **Figure 119 - Segmentation and Reassembly Headers**

3342 **Table 58 - 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: Last segment with SAR.CRC field at the end of the segment.
SAR.NSEGS	6 bits	‘Number of Segments’ – 1. Note: This field is only present in segments with SAR.TYPE=0b00
SAR.SEQ	6 bits	Sequence number of segment. Note: This field is only present in segments with SAR.TYPE=0b01, SAR.TYPE=0b10 and SAR.TYPE=0b11.

3343 Every segment (except for the first one) includes a sequence number so that the loss of a segment could be
 3344 detected in reassembly. The sequence numbering shall start from zero with every new CPCS SDU. The first
 3345 segment which contains a SAR.SEQ field must have SAR.SEQ = 0. All subsequent segments from the same
 3346 CPCS SDU shall increase this sequence number such that the SAR.SEQ field adds one with every transmission.

3347 The value SAR.NSEGS indicates the total number of segments, minus one. So when SAR.NSEGS = 0, the CPCS
 3348 SDU is sent in one segment. SAR.NSEGS = 63 indicates there will be 64 segments to form the full CPCS SDU.
 3349 When SAR.NSEGS = 0, it indicates that this first segment is also the last segment. No further segment with
 3350 SAR.TYPE = 0b01 or 0b10 is to be expected for this one-segment CPCS SDU.

3351 Using segments with SAR.TYPE=0b11 instead of SAR.TYPE=0b10 will be recommended for the last segments
 3352 of connections without ARQ, multicast and broadcast. Connections with ARQ may use SAR.TYPE=0b10
 3353 safely (this reduces overhead and guarantees backward compatibility). The 32 bits CRC is computed using
 3354 the polynomial generator in the and appended at the end of the last segment.

3355 **Table 59 - SAR.CRC**

Name	Length	Description
SAR.CRC	32 bits	CRC32 of the SAR CPCS PDU. This field is only present in segments with SAR.TYPE=0b11. 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.

3356 SAR at the receiving end shall buffer all segments and deliver only fully reassembled CPCS SDUs to the SSCS
 3357 above. Should reassembly fail due to a segment not being received or too many segments being ...received
 3358 etc., SAR shall not deliver any incomplete CPCS SDU to the SSCS above.

3359 **5.2.2.2 SAR constants**

3360 Table 53 shows the constants for the SAR service.

3361 **Table 60 - SAR Constants**

Constant	Value
CIMaxAppPktSize	Max Value (SAR.NSEGS) x macSARSize.

3362 5.3 NULL Service-Specific Convergence Sublayer (NULL SSCS)

3363 5.3.1 Overview

3364 Null SSCS provides the MAC layer with a transparent path to upper layers, being as simple as possible and
3365 minimizing overhead. It is intended for applications that do not need any special convergence capability.

3366 The unicast and multicast connections of this SSCS shall use the SAR service, as defined in 5.2.2. If they do
3367 not need the SAR service they shall still include the SAR header (notifying just one segment).

3368 The CON.TYPE and MUL.TYPE (see Annex E) for unicast connections and multicast groups shall use the same
3369 type that has been already defined for the application that makes use of this Null SSCS.

3370 5.3.2 Primitives

3371 Null SSCS primitives are just a direct mapping of the MAC primitives. A full description of every primitive is
3372 avoided, because the mapping is direct and they will work as the ones of the MAC layer.

3373 The directly mapped primitives have exactly the same parameters as the ones in the MAC layer and perform
3374 the same functionality. The set of primitives that are directly mapped are shown below.

3375 **Table 61 - 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

Null SCS mapped to a MAC primitive
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

3376 5.4 IPv4 Service-Specific Convergence Sublayer (IPv4 SCS)

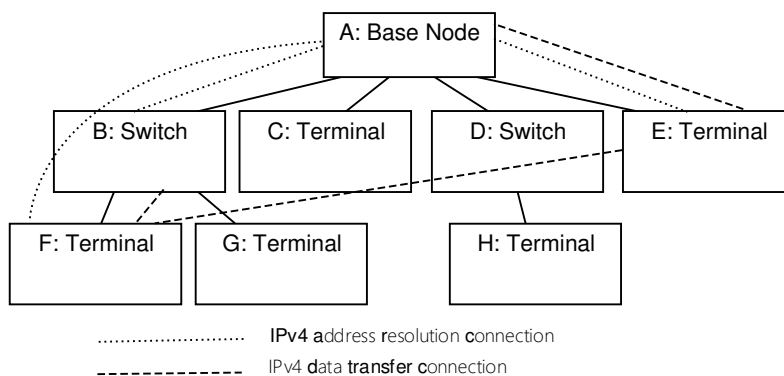
3377 5.4.1 Overview

3378 The IPv4 SCS provides an efficient method for transferring IPv4 packets over the PRIME Subnetworks.
3379 Several conventions do apply:

- 3380 • A Service Node can send IPv4 packets to the Base Node or to other Service Nodes.
- 3381 • It is assumed that the Base Node acts as a router between the PRIME Subnetwork and any other
3382 network. The Base Node could also act as a NAT. How the Base Node connects to the other
3383 networks is beyond the scope of this specification.
- 3384 • In order to keep implementations simple, only one single route is supported per local IPv4
3385 address.
- 3386 • Service Nodes may use statically configured IPv4 addresses or DHCP to obtain IPv4 addresses.
- 3387 • The Base Node performs IPv4 to EUI-48 address resolution. Each Service Node registers its IPv4
3388 address and EUI-48 address with the Base Node (see section 5.4.2). Other Service Nodes can then
3389 query the Base Node to resolve an IPv4 address into a EUI-48 address. This requires the
3390 establishment of a dedicated connection with the Base Node for address resolution.
- 3391 • The IPv4 SCS performs the routing of IPv4 packets. In other words, the IPv4 SCS will decide
3392 whether the packet should be sent directly to another Service Node or forwarded to the
3393 configured gateway.
- 3394 • Although IPv4 is a connectionless protocol, the IPv4 SCS is connection-oriented. Once address
3395 resolution has been performed, a connection is established between the source and destination
3396 Service Node for the transfer of IPv4 packets. This connection is maintained while traffic is being
3397 transferred and may be closed after a period of inactivity.
- 3398 • The CPCS (see section 5.2) SAR sublayer shall always be present with the IPv4 Convergence layer.
3399 Generated MSDUs are at most 'macSARSize' bytes long and upper layer PDU messages are not
3400 expected must not to be longer than CIMaxAppPktSize.

- 3401 • Optionally TCP/IPv4 headers may be compressed. Compression is negotiated as part of the
- 3402 connection establishment phase.
- 3403 • The broadcasting of IPv4 packets is supported using the MAC broadcast mechanism.
- 3404 • The multicasting of IPv4 packets is supported using the MAC multicast mechanism.

3405 The IPv4 SCS has a number of connection types. For address resolution there is a connection to the Base
 3406 Node. For IPv4 data transfer there is one connection per Destination Node: with the Base Node that acts as
 3407 the IPv4 gateway to other networks or to/with any other Node in the same Subnetwork. This is shown in
 3408 Figure 120.



3409
 3410 **Figure 120 - IPv4 SCS connection example**

3411 Here, Nodes B, E and F have address resolution connections to the Base Node. Node E has a data connection
 3412 to the Base Node and Node F. Node F is also has a data connection to Node B. The figure does not show
 3413 broadcast and multicast connections.

3414 5.4.2 Address resolution

3415 5.4.2.1 General

3416 The IPv4 layer will present the IPV4 SCS with an IPV4 packet to be transferred. The IPV4 SCS is responsible
 3417 for determining which Service Node the packet should be delivered to using the IPV4 addresses in the packet.
 3418 The IPV4 SCS must then establish a connection to the destination if one does not already exist so that the
 3419 packet can be transferred. Three classes of IPV4 addresses can be used and the following subsections describe
 3420 how these addresses are resolved into EUI-48 addresses.

3421 5.4.2.2 Unicast addresses

3422 5.4.2.2.1 General

3423 IPv4 unicast addresses must be resolved into unicast EUI-48 addresses. The Base Node maintains a database
 3424 of IPV4 addresses and EUI-48 addresses. Address resolution then operates by querying this database. A
 3425 Service Node must establish a connection to the address resolution service running on the Base Node, using
 3426 the connection type value TYPE (see Annex E) TYPE_CL_IPV4_AR. No data should be passed in the connection

3427 establishment. Using this connection, the Service Node can use two mechanisms as defined in the following
3428 paragraphs.

3429 **5.4.2.2.2 Address registration and unregistration**

3430 A Service Node uses the AR_REGISTER_S message to register an IPv4 address and the corresponding EUI-48
3431 address meaning request from the base node to record inside its registration table, the IPv4 address and its
3432 corresponding service node EUI-48. The Base Node will acknowledge an AR_REGISTER_B message. The
3433 Service Node may register multiple IPv4 addresses for the same EUI-48 address.

3434 A Service Node uses the AR_DEREGISTER_S message to unregister an IPv4 address and the corresponding
3435 EUI-48 address meaning requests from the base node to delete inside its registration table, the entry
3436 corresponding to the concerned IPv4 address. The Base Node will acknowledge it with an AR_DEREGISTER_B
3437 message.

3438 When the IPv4 address resolution connection between the Service Node and the Base Node is closed, the
3439 Base Node should remove all addresses associated to that connection.

3440 **5.4.2.2.3 Address lookup**

3441 A Service Node uses the AR_LOOKUP_S message to perform a lookup. The message contains the IPv4 address
3442 to be resolved. The Base Node will respond with an AR_LOOKUP_B message that contains an error code and,
3443 if there is no error, the EUI-48 address associated with the IPv4 address. If the Base Node has multiple entries
3444 in its database for the same IPv4 address, the possible returned EUI-48 address is undefined.

3445 **5.4.2.3 Broadcast Address**

3446 IPv4 broadcast address 255.255.255.255 maps to a MAC broadcast connection with LCID equal to
3447 LCI_CL_IPv4_BROADCAST. All IPv4 broadcast packets will be sent to this connection. When an IPv4 broadcast
3448 packet is received on this connection, the IPv4 address should be examined to determine if it is a broadcast
3449 packet for the Subnetwork in which the Node has an IPv4 address. Only broadcast packets from member
3450 subnets should be passed up the IPv4 protocol stack.

3451 **5.4.2.4 Multicast Addresses**

3452 Multicast IPv4 addresses are mapped to a PRIME MAC multicast connection by the Base Node using an
3453 address resolution protocol.

3454 To join a multicast group, AR_MCAST_REG_S is sent from the Service Node to the Base Node with the IPv4
3455 multicast address. The Base Node will reply with an AR_MCAST_REG_B that contains the LCID value assigned
3456 to the said multicast address. However, the Base Node may also allocate other LCIDs which are not in use if
3457 it so wishes. The Service Node can then join a multicast group (see 4.6.7.2) for the given LCID to receive IPv4
3458 multicast packets. These LCID values can be reused so that multiple IPv4 destination multicast addresses can
3459 be seen on the same LCID. To leave the multicast group, AR_MCAST_UNREG_S is sent from the Service Node
3460 to the Base Node with the IPv4 multicast address. The Base Node will acknowledge it with an
3461 AR_MCAST_UNREG_B message.

3462 When a Service Node wants to send an IPv4 multicast datagram, it just uses the appropriate LCID. If the
 3463 Service Node has not joined the multicast group, it needs first to learn the LCID to be used. The process with
 3464 AR_MCAST_REG_{S|B} messages as described above can be used. While IPv4 multicast packets are still being
 3465 sent, the Service Node remains registered to the multicast group. T_{mcast_reg} after the last IPv4 multicast
 3466 datagram was sent, the Service Node should unregister from the multicast group, by means of
 3467 AR_MCAST_UNREG_{S|B} messages. The nominal value of T_{mcast_reg} is 10 minutes; however, other values
 3468 may be used.

3469 5.4.2.5 Retransmission of address resolution packets

3470 The connection between the Service Node and the Base Node for address resolution is not reliable if the MAC
 3471 ARQ is not used. The Service Node is responsible for making retransmissions if the Base Node does not
 3472 respond in one second. It is not considered an error when the Base Node receives the same registration
 3473 requests multiple times or is asked to remove a registration that does not exist. These conditions can be the
 3474 result of retransmissions.

3475 5.4.3 IPv4 packet transfer

3476 For packets to be transferred, a connection needs to be established between source and Destination Nodes.
 3477 The IPV4 SSCS will examine each IPv4 packet to determine the destination EUI-48 address. If a data
 3478 connection to the destination already exists, the packet is sent. To establish this, IPv4 SSCS keeps a table for
 3479 each connection, with information shown in Table 55 (see RFC 1144).. To use this table, it is first necessary
 3480 to determine if the IPv4 destination address is in the local Subnetwork or if a gateway has to be used. The
 3481 netmask associated with the local IPv4 address is used to determine this. If the IPv4 destination address is
 3482 not in the local Subnetwork, the address of the default gateway is used instead of the destination address
 3483 when the table is searched.

3484 **Table 62 - 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.

3485 The IPV4 SSCS may close a connection when it has not been used for an implementation-defined time period.
3486 When the connection is closed the entry for the connection is removed at both ends of the connection.

3487 When a connection to the destination does not exist, more work is necessary. The address resolution service
3488 is used to determine the EUI-48 address of the remote IPv4 address if it is local or the gateway associated
3489 with the local address if the destination address is in another Subnetwork. When the Base Node replies with
3490 the EUI-48 address of the destination Service Node, a MAC connection is established to the remote device.
3491 The TYPE value of this connection is TYPE_CL_IPv4_UNICAST. The data passed in the request message is
3492 defined in section 5.4.7.4. The local IPv4 address is provided so that the remote device can add the new
3493 connection to its cache of connections for sending data in the opposite direction. The use of Van Jacobson
3494 Header Compression is also negotiated as part of the connection establishment. Once the connection has
3495 been established, the IPv4 packet can be sent.

3496 When the packet is addressed to the IPv4 broadcast address, the packet has to be sent using the MAC
3497 broadcast service. When the IPV4 SSCS is opened, a broadcast connection is established for transferring all
3498 broadcast packets. The broadcast IPv4 packet is simply sent to this connection. Any packet received on this
3499 broadcast connection is passed to the IPv4 protocol stack.

3500 **5.4.4 Segmentation and reassembly**

3501 The IPV4 SSCS should support IPv4 packets with an MTU of 1500 bytes. This requires the use of SAR (see
3502 5.2.2).

3503 **5.4.5 Header compression**

3504 Van Jacobson TCP/IP Header Compression is an optional feature in the IPv4 SSCS. The use of VJ compression
3505 is negotiated as part of the connection establishment phase of the connection between two Service Nodes.

3506 VJ compression is designed for use over a point-to-point link layer that can inform the decompressor when
3507 packets have been corrupted or lost. When there are errors or lost packets, the decompressor can then
3508 resynchronize with the compressor. Without this resynchronization process, erroneous packets will be
3509 produced and passed up the IPv4 stack.

3510 The MAC layer does not provide the facility of detecting lost packets or reporting corrupt packets. Thus, it is
3511 necessary to add this functionality in the IPV4 SSCS. The IPV4 SSCS maintains two sequence numbers when
3512 VJ compression is enabled for a connection. These sequence numbers are 8 bits in size. When transmitting
3513 an IPv4 packet, the CL_IPv4_CON.TxSeq sequence number is placed in the packet header, as shown in Section
3514 5.4.3. The sequence number is then incremented. Upon reception of a packet, the sequence number in the
3515 received packet is compared against CL_IPv4_CON.RxSeq. If they differ, TYPE_ERROR, as defined in RFC1144,
3516 is passed to the decompressor. The CL_IPv4_CON.RxSeq value is always updated to the value received in the
3517 packet header.

3518 Header compression should never be negotiated for broadcast or multicast packets.

3519 5.4.6 Quality of Service mapping

3520 The PRIME MAC specifies that the contention-based access mechanism supports 4 priority levels (1-4). Level
3521 1 is used for MAC control messages, but not exclusively so.

3522 IPv4 packets include a TOS field in the header to indicate the QoS the packet would like to receive. Three bits
3523 of the TOS indicate the IP Precedence. The following table specifies how the IP Precedence is mapped into
3524 the PRIME MAC priority.

3525 **Table 63 - Mapping IPv4 Precedence to PRIME MAC priority**

IP Precedence	MAC Priority
000 – Routine	3
001 – Priority	3
010 – Immediate	2
011 – Flash	2
100 – Flash Override	1
101 – Critical	1
110 – Internetwork Control	0
111 – Network Control	0

3526 5.4.7 Packet formats and connection data

3527 5.4.7.1 General

3528 This section defines the format of IPV4 SSCS PDUs.

3529 **5.4.7.2 Address resolution PDUs**

3530 **5.4.7.2.1 General**

3531 The following PDUs are transferred over the address resolution connection between the Service Node and
 3532 the Base Node. The following sections define AR.MSG values in the range of 0 to 11. All higher values are
 3533 reserved for later versions of this specification.

3534 **5.4.7.2.2 AR_REGISTER_S**

3535 Table 64 shows the address resolution register message sent from the Service Node to the Base Node.

3536 **Table 64 - AR_REGISTER_S message format**

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type. <ul style="list-style-type: none"> For AR_REGISTER_S = 0.
AR.IPv4	32-bits	IPv4 address to be registered.
AR.EUI-48	48-bits	EUI-48 to be registered.

3537 **5.4.7.2.3 AR_REGISTER_B**

3538 Table 65 shows the address resolution register acknowledgment message sent from the Base Node to the
 3539 Service Node.

3540 **Table 65 - AR_REGISTER_B message format**

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type. <ul style="list-style-type: none"> For AR_REGISTER_B = 1.
AR.IPv4	32-bits	Registered IPv4 address.
AR.EUI-48	48-bits	EUI-48 registered.

3541

3542 The AR.IPv4 and AR.EUI-48 fields are included in the AR_REGISTER_B message so that the Service Node can
3543 perform multiple overlapping registrations.

3544 5.4.7.2.4 AR_UNREGISTER_S

3545 Table 66 shows the address resolution unregister message sent from the Service Node to the Base Node.

3546 **Table 66 - AR_UNREGISTER_S message format**

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type. <ul style="list-style-type: none"> For AR_UNREGISTER_S = 2.
AR.IPv4	32-bits	IPv4 address to be unregistered.
AR.EUI-48	48-bits	EUI-48 to be unregistered.

3547 5.4.7.2.5 AR_UNREGISTER_B

3548 Table 67 shows the address resolution unregister acknowledgment message sent from the Base Node to the
3549 Service Node.

3550 **Table 67 - AR_UNREGISTER_B message format**

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type. <ul style="list-style-type: none"> For AR_UNREGISTER_B = 3.
AR.IPv4	32-bits	Unregistered IPv4 address .
AR.EUI-48	48-bits	Unregistered EUI-48.

3551 The AR.IPv4 and AR.EUI-48 fields are included in the AR_UNREGISTER_B message so that the Service Node
3552 can perform multiple overlapping Unregistrations.

3553 5.4.7.2.6 AR_LOOKUP_S

3554 Table 68 shows the address resolution lookup message sent from the Service Node to the Base Node.

3555

Table 68 - AR_LOOKUP_S message format

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type. <ul style="list-style-type: none"> For AR_LOOKUP_S = 4.
AR.IPv4	32-bits	IPv4 address to lookup.

3556 **5.4.7.2.7 AR_LOOKUP_B**

3557 Table 69 shows the address resolution lookup response message sent from the Base Node to the Service
3558 Node.

3559

Table 69 - 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	Lookup 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.

3560 The lookup may fail if the requested address has not been registered. In that case, AR.Status will have a value
3561 other than zero and the contents of AR.EUI-48 will be undefined. The lookup is only successful when
3562 AR.Status is zero. In that case, the EUI-48 field contains the resolved address.

3563 **5.4.7.2.8 AR_MCAST_REG_S**

3564 Table 70 shows the multicast address resolution register message sent from the Service Node to the Base
3565 Node.

3566

Table 70 - 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.

3567 **5.4.7.2.9 AR_MCAST_REG_B**

3568 Table 71 shows the multicast address resolution register acknowledgment message sent from the Base Node
3569 to the Service Node.

3570

Table 71 - 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.
<i>Reserved</i>	2-bits	Reserved. Should be encoded as 0.
AR.LCID	6-bits	LCID assigned to this IPv4 multicast address.

3571 The AR.IPv4 field is included in the AR_MCAST_REG_B message so that the Service Node can perform multiple
3572 overlapping registrations.

3573 **5.4.7.2.10 AR_MCAST_UNREG_S**

3574 Table 72 shows the multicast address resolution unregister message sent from the Service Node to the Base
3575 Node.

3576

Table 72 - AR_MCAST_UNREG_S message format

Name	Length	Description
------	--------	-------------

AR.MSG	8-bits	Address Resolution Message Type. • For AR_MCAST_UNREG_S = 10.
AR.IPv4	32-bits	IPv4 multicast address to be unregistered.

3577 **5.4.7.2.11 AR_MCAST_UNREG_B**

3578 Table 73 shows the multicast address resolution unregister acknowledgment message sent from the Base
3579 Node to the Service Node.

3580 **Table 73 - AR_MCAST_UNREG_B message format**

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type. • For AR_MCAST_UNREG_B = 11;
AR.IPv4	32-bits	IPv4 multicast address unregistered.

3581 The AR.IPv4 field is included in the AR_MCAST_UNREG_B message so that the Service Node can perform
3582 multiple overlapping Unregistrations.

3583 **5.4.7.3 IPv4 packet format**

3584 **5.4.7.3.1 General**

3585 The following PDU formats are used for transferring IPv4 packets between Service Nodes. Two formats are
3586 defined. The first format is for when header compression is not used. The second format is for Van Jacobson
3587 Header Compression.

3588 **5.4.7.3.2 IPv4 Packet Format, No Negotiated Header Compression**

3589 When no header compression has been negotiated, the IPv4 packet is simply sent as is, without any header.

3590 **Table 74 - IPv4 Packet format without negotiated header compression**

Name	Length	Description
<i>IPv4.PKT</i>	n-octets	The IPv4 Packet.

3591 **5.4.7.3.3 IPv4 Packet Format with VJ Header Compression**

3592 With Van Jacobsen header compression, a one-octet header is needed before the IPv4 packet.

3593 **Table 75 - 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.

3594 The IPv4.Type value TYPE_ERROR is never sent. It is a pseudo packet type used to tell the decompressor that
3595 a packet has been lost.3596 **5.4.7.4 Connection Data**3597 **5.4.7.4.1 General**3598 When a connection is established between Service Nodes for the transfer of IPv4 packets, data is also
3599 transferred in the connection request packets. This data allows the negotiation of compression and
3600 notification of the IPv4 address.3601 **5.4.7.4.2 Connection Data from the Initiator**

3602 Table 76 shows the connection data sent by the initiator.

3603 **Table 76 - Connection data sent by the initiator**

Name	Length	Description
<i>Reserved</i>	6-bits	Should be encoded as 0 in this version of the IPV4 SCS protocol.
Data.HC	2-bit	Header Compression . <ul style="list-style-type: none"> • Data.HC = 0 – No compression requested;

Name	Length	Description
		<ul style="list-style-type: none"> • Data.HC = 1 – VJ Compression requested; • Data.HC = 2, 3 – Reserved for future versions of the specification.
Data.IPv4	32-bits	IPv4 address of the initiator

3604 If the device accepts the connection, it should copy the Data.IPv4 address into a new table entry along with
 3605 the negotiated Data.HC value.

3606 5.4.7.4.3 Connection Data from the Responder

3607 Table 77 shows the connection data sent in response to the connection request.

3608 **Table 77 - Connection data sent by the responder**

Name	Length	Description
<i>Reserved</i>	6-bits	Should be encoded as zero in this version of the IPV4 SCS 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.

3609 A header compression scheme can only be used when it is supported by both Service Nodes. The responder
 3610 may only set Data.HC to 0 or the same value as that received from the initiator. When the same value is used,
 3611 it indicates that the requested compression scheme has been negotiated and will be used for the connection.
 3612 Setting Data.HC to 0 allows the responder to deny the request for that header compression scheme or force
 3613 the use of no header compression.

3614 5.4.8 Service Access Point

3615 5.4.8.1 General

3616 This section defines the service access point used by the IPv4 layer to communicate with the IPV4 SCS.

3617 **5.4.8.2 Opening and closing the IPv4 SSCS**

3618 **5.4.8.2.1 General**

3619 The following primitives are used to open and close the IPv4 SSCS. The IPv4 SSCS may be opened once only.
3620 The IPv4 layer may close the IPv4 SSCS when the IPv4 interface is brought down. The IPv4 SSCS will also close
3621 the IPv4 SSCS when the underlying MAC connection to the Base Node has been lost.

3622 **5.4.8.2.2 CL_IPv4_ESTABLISH.request**

3623 The CL_IPv4_ESTABLISH.request primitive is passed from the IPv4 layer to the IPV4 SSCS. It is used when the
3624 IPv4 layer brings the interface up.

3625 The semantics of this primitive are as follows:

3626 *CL_IPv4_ESTABLISH.request{AE}*

3627 The AE parameter indicates whether the interface will be authenticated and encrypted or not.

3628 On receiving this primitive, the IPV4 SSCS will form the address resolution connection to the Base Node and
3629 join the broadcast group used for receiving/transmitting broadcast packets.

3630 **5.4.8.2.3 CL_IPv4_ESTABLISH.confirm**

3631 The CL_IPv4_ESTABLISH.confirm primitive is passed from the IPV4 SSCS to the IPv4 layer. It is used to indicate
3632 that the IPV4 SSCS is ready to access IPv4 packets to be sent to peers.

3633 The semantics of this primitive are as follows:

3634 *CL_IPv4_ESTABLISH.confirm{AE}*

3635 The AE parameter indicates whether the interface will be authenticated and encrypted or not.

3636 Once the IPv4 SSCS has established all the necessary connections and is ready to transmit and receive IPv4
3637 packets, this primitive is passed to the IPv4 layer. If the IPV4 SSCS encounters an error while opening, it
3638 responds with a CL_IPv4_RELEASE.confirm primitive, rather than a CL_IPv4_ESTABLISH.confirm.

3639 **5.4.8.2.4 CL_IPv4_RELEASE.request**

3640 The CL_IPv4_RELEASE.request primitive is used by the IPv4 layer when the interface is put down. The IPV4
3641 SSCS closes all connections so that no more IPv4 packets are received and all resources are released.

3642 The semantics of this primitive are as follows:

3643 *CL_IPv4_RELEASE.request{}*

3644 Once the IPV4 SSCS has released all its connections and resources it returns a CL_IPv4_RELEASE.confirm.

3645 **5.4.8.2.5 CL_IPv4_RELEASE.confirm**

3646 The CL_IPv4_RELEASE.confirm primitive is used by the IPv4 SCS to indicate to the IPv4 layer that the IPv4
3647 SCS has been closed. This can be as a result of a CL_IPv4_RELEASE.request primitive, a
3648 CL_IPv4_ESTABLISH.request primitive, or because the MAC layer indicates the address resolution connection
3649 has been lost, or the Service Node itself is no longer registered.

3650 The semantics of this primitive are as follows:

3651 *CL_IPv4_RELEASE.confirm{result}*

3652 The result parameter has the meanings defined in Table 142.

3653 **5.4.8.3 Unicast address management**

3654 **5.4.8.3.1 General**

3655 The primitives defined here are used for address management, i.e. the registration and Unregistration of IPv4
3656 addresses associated with this IPv4 SCS .

3657 When there are no IPv4 addresses associated with the IPv4 SCS, the IPv4 SCS will only send and receive
3658 broadcast and multicast packets; unicast packets may not be sent. However, this is sufficient for
3659 BOOTP/DHCP operation to allow the device to gain an IPv4 address. Once an IPv4 address has been
3660 registered, the IPv4 layer can transmit unicast packets that have a source address equal to one of its
3661 registered addresses.

3662 **5.4.8.3.2 CL_IPv4_REGISTER.request**

3663 This primitive is passed from the IPv4 layer to the IPv4 SCS to register an IPv4 address.

3664 The semantics of this primitive are as follows:

3665 *CL_IPv4_REGISTER.request{IPv4, netmask, gateway}*

3666 The IPv4 address is the address to be registered.

3667 The netmask is the network mask, used to mask the network number from the address. The netmask is used
3668 by the IPv4 SCS to determine whether the packet should be delivered directly or the gateway should be
3669 used.

3670 The gateway is an IPv4 address of the gateway to be used for packets with the IPv4 local address but the
3671 destination address is not in the same Subnetwork as the local address.

3672 Once the IPv4 address has been registered to the Base Node, a CL_IPv4_REGISTER.confirm primitive is used.
3673 If the registration fails, the CL_IPv4_RELEASE.confirm primitive will be used.

3674 **5.4.8.3.3 CL_IPv4_REGISTER.confirm**

3675 This primitive is passed from the IPv4 SSCS to the IPv4 layer to indicate that a registration has been successful.

3676 The semantics of this primitive are as follows:

3677 *CL_IPv4_REGISTER.confirm{IPv4}*

3678 The IPv4 address is the address that was registered.

3679 Once registration has been completed, the IPv4 layer may send IPv4 packets using this source address.

3680 **5.4.8.3.4 CL_IPv4_UNREGISTER.request**

3681 This primitive is passed from the IPv4 layer to the IPv4 SSCS to unregister an IPv4 address.

3682 The semantics of this primitive are as follows:

3683 *CL_IPv4_UNREGISTER.request{IPv4}*

3684 The IPv4 address is the address to be unregistered.

3685 Once the IPv4 address has been unregistered to the Base Node, a CL_IPv4_UNREGISTER.confirm primitive is
3686 used. If the unregistration fails, the CL_IPv4_RELEASE.confirm primitive will be used.

3687 **5.4.8.3.5 CL_IPv4_UNREGISTER.confirm**

3688 This primitive is passed from the IPv4 SSCS to the IPv4 layer to indicate that an Unregistration has been
3689 successful.

3690 The semantics of this primitive are as follows:

3691 *CL_IPv4_UNREGISTER.confirm{IPv4}*

3692 The IPv4 address is the address that was unregistered.

3693 Once Unregistration has been completed, the IPv4 layer may not send IPv4 packets using this source address.

3694 **5.4.8.4 Multicast group management**

3695 **5.4.8.4.1 General**

3696 This section describes the primitives used to manage multicast groups.

3697 **5.4.8.4.2 CL_IPv4_IGMP_JOIN.request**

3698 This primitive is passed from the IPv4 layer to the IPv4 SSCS. It contains an IPv4 multicast address that is to
3699 be joined.

3700 The semantics of this primitive are as follows:

3701 *CL_IPv4_IGMP_JOIN.request{IPv4, AE }*

3702 The IPv4 address is the IPv4 multicast group that is to be joined.

3703 The AE parameter indicates whether messages in this group will be authenticated and encrypted or not.

3704 When the IPv4 SSCS receives this primitive, it will arrange for IPv4 packets sent to this group to be multicast
3705 in the PRIME network and receive packets using this address to be passed to the IPv4 stack. If the IPv4 SSCS
3706 cannot join the group, it uses the *CL_IPv4_IGMP_LEAVE.confirm* primitive. Otherwise the
3707 *CL_IPv4_IGMP_JOIN.confirm* primitive is used to indicate success.

3708 **5.4.8.4.3 CL_IPv4_IGMP_JOIN.confirm**

3709 This primitive is passed from the IPv4 SSCS to the IPv4. It contains a result status and an IPv4 multicast address
3710 that was joined.

3711 The semantics of this primitive are as follows:

3712 *CL_IPv4_IGMP_JOIN.confirm{IPv4, AE }*

3713 The IPv4 address is the IPv4 multicast group that was joined. The IPv4 SSCS will start forwarding IPv4 multicast
3714 packets for the given multicast group.

3715 The AE parameter indicates whether the messages in this group will be authenticated and encrypted or not.

3716 **5.4.8.4.4 CL_IPv4_IGMP_LEAVE.request**

3717 This primitive is passed from the IPv4 layer to the IPv4 SSCS. It contains an IPv4 multicast address to be left.

3718 The semantics of this primitive are as follows:

3719 *CL_IPv4_IGMP_LEAVE.request{IPv4}*

3720 The IPv4 address is the IPv4 multicast group to be left. The IPv4 SSCS will stop forwarding IPv4 multicast
3721 packets for this group and may leave the PRIME MAC multicast group.

3722 **5.4.8.4.5 CL_IPv4_IGMP_LEAVE.confirm**

3723 This primitive is passed from the IPv4 SSCS to the IPv4. It contains a result status and an IPv4 multicast address
3724 that was left.

3725 The semantics of this primitive are as follows:

3726 *CL_IPv4_IGMP_LEAVE.confirm{IPv4, Result}*

3727 The IPv4 address is the IPv4 multicast group that was left. The IPv4 SSCS will stop forwarding IPv4 multicast
3728 packets for the given multicast group.

3729 The Result takes a value from Table 142.

3730 This primitive can be used by the IPv4 SSCS as a result of a CL_IPv4_IGMP_JOIN.request,
3731 CL_IPv4_IGMP_LEAVE.request or because of an error condition resulting in the loss of the PRIME MAC
3732 multicast connection.

3733 **5.4.8.5 Data transfer**

3734 **5.4.8.5.1 General**

3735 The following primitives are used to send and receive IPv4 packets.

3736 **5.4.8.5.2 CL_IPv4_DATA.request**

3737 This primitive is passed from the IPv4 layer to the IPv4 SSCS. It contains one IPv4 packet to be sent.

3738 The semantics of this primitive are as follows:

3739 *CL_IPv4_DATA.request{IPv4_PDU}*

3740 The IPv4_PDU is the IPv4 packet to be sent.

3741 **5.4.8.5.3 CL_IPv4_DATA.confirm**

3742 This primitive is passed from the IPv4 SSCS to the IPv4 layer. It contains a status indication and an IPv4 packet
3743 that has just been sent.

3744 The semantics of this primitive are as follows:

3745 *CL_IPv4_DATA.confirm{IPv4_PDU, Result}*

3746 The IPv4_PDU is the IPv4 packet that was to be sent.

3747 The Result value indicates whether the packet was sent or an error occurred. It takes a value from Table 142.

3748 **5.4.8.5.4 CL_IPv4_DATA.indicate**

3749 This primitive is passed from the IPv4 SSCS to the IPv4 layer. It contains an IPv4 packet that has just been
3750 received.

3751 The semantics of this primitive are as follows:

3752 *CL_IPv4_DATA.indicate{IPv4_PDU }*

3753 The IPv4_PDU is the IPv4 packet that was received.

3754 **5.5 IEC 61334-4-32 Service-Specific Convergence Sublayer (IEC** 3755 **61334-4-32 SCS)**

3756 **5.5.1 General**

3757 For all the service required, the IEC 61334-4-32 SCS supports the DL_DATA primitives as defined in the IEC
3758 61334-4-32 standard. IEC 61334-4-32 should be read at the same time as this section, which is not standalone
3759 text.

3760 **5.5.2 Overview**

3761 The IEC 61334-4-32 SCS provides convergence functions for applications that use IEC 61334-4-32 services.
3762 Implementations conforming to this SCS shall offer all LLC basic and management services as specified in
3763 IEC 61334-4-32 (1996-09 Edition), subsections 2.2.1 and 2.2.3. Additionally, the IEC 61334-4-32 SCS specified
3764 in this section provides extra services that help mapping this connection-less IEC 61334-4-32 LLC protocol to
3765 the connection-oriented nature of MAC.

- 3766 • A Service Node can only exchange data with the Base Node and not with other Service Nodes.
3767 This meets all the requirements of IEC 61334-4-32, which has similar restrictions.
- 3768 • Each IEC 61334-4-32 SCS session establishes a dedicated PRIME MAC connection for exchanging
3769 unicast data with the Base Node.
- 3770 • The Service Node SCS session is responsible for initiating this connection to the Base Node. The
3771 Base Node SCS cannot initiate a connection to a Service Node.
- 3772 • Each IEC 61334-4-32 SCS listens to a PRIME broadcast MAC connection dedicated to the transfer
3773 of IEC 61334-4-32 broadcast data from the Base Node to the Service Nodes. This broadcast
3774 connection is used when applications in the Base Node using IEC 61334-4-32 services make a
3775 transmission request with the Destination_address used for broadcast or the broadcast SAP
3776 functions are used. When there are multiple SCS sessions within a Service Node, one PRIME
3777 broadcast MAC connection is shared by all the SCS sessions.
- 3778 • A CPCS session is always present with a IEC 61334-4-32 SCS session. The SPCS sublayer
3779 functionality is as specified in Section 5.2.2. Thus, the MSDUs generated by IEC 61334-4-32 SCS
3780 are always less than *macSARSize* bytes and application messages shall not be longer than
3781 *CIMaxAppPktSize*.

3782 **5.5.3 Address allocation and connection establishment**

3783 Each 4-32 connection will be identified with the "Application unique identifier" that will be communicating
3784 through this 4-32 connection. It is the scope of the communication profile based on these lower layers to
3785 define the nature and rules for, this unique identifier. Please refer to the future prTS/EN52056-8-4 for the
3786 DLMS/COSEM profile unique identifier. As long as the specification of the 4-32 Convergence layer concerns
3787 this identifier will be called the "Device Identifier".

3788 The protocol stack as defined in IEC 61334 defines a Destination address to identify each device in the
3789 network. This Destination address is specified beyond the scope of the IEC 61334-4-32 document. However,
3790 it is used by the document. So that PRIME devices can make use of the 4-32 layer, this Destination address is
3791 also required and is specified here. For more information about this Destination address, please see IEC
3792 61334-4-1 section 4.3, MAC Addresses.

3793 The Destination address has a scope of one PRIME Subnetwork. The Base Node 4-32 SSCP layer is responsible
3794 for allocating these addresses dynamically and associating the Device Identifier of the Service Nodes SSCP
3795 session device with the allocated Destination address, according to the IEC-61334-4-1 standard. The
3796 procedure is as follows:

3797 When the Service Node IEC 61334-4-32 SSCP session is opened by the application layer, it passes the Device
3798 Identifier of the device. The IEC 61334-4-32 SSCP session then establishes its unicast connection to the Base
3799 Node. This unicast connection uses the PRIME MAC TYPE value TYPE_CL_432, as defined in Table 140. The
3800 connection request packet sent from the Service Node to the Base Node contains a data parameter. This data
3801 parameter contains the Device Identifier. The format of this data is specified in section 5.5.4.2.

3802 On receiving this connection request at the Base Node, the Base Node allocates a unique Subnetwork
3803 Destination address to the Service Nodes SSCP session. The Base Node sends back a PRIME MAC connection
3804 response packet that contains a data parameter. This data parameter contains the allocated Destination
3805 address and the address being used by the Base Node itself. The format of this data parameter is defined in
3806 section 5.5.4.2. A 4-32 CL SAP primitive is used in the Base Node to indicate this new Service Node SSCP
3807 session mapping of Device Identifier and Destination_address to the 4-32 application running in the Base
3808 Node.

3809 On receiving the connection establishment and the Destination_address passed in the PRIME MAC
3810 connection establishment packet, the 4-32 SSCP session confirms to the application that the Convergence
3811 layer session has been opened and indicates the Destination_address allocated to the Service Node SSCP
3812 session and the address of the Base Node. The Service Node also opens a PRIME MAC broadcast connection
3813 with LCID equal to LCI_CL_432_BROADCAST, as defined in Table 141, if no other SSCP session has already
3814 opened such a broadcast connection. This connection is used to receive broadcast packets sent by the Base
3815 Node 4-32 Convergence layer to all Service Node 4-32 Convergence layer sessions.

3816 If the Base Node has allocated all its available Destination_addresses, due to the exhaustion of the address
3817 space or implementation limits, it should simply reject the connection request from the Service Node. The
3818 Service Node may try to establish the connection again. However, to avoid overloading the PRIME
3819 Subnetwork with such requests, it should limit such connection establishments to one attempt per minute
3820 when the Base Node rejects a connection establishment.

3821 When the unicast connection between a Service Node and the Base Node is closed (e.g. because the
3822 Convergence layer on the Service Node is closed or the PRIME MAC level connection between the Service
3823 Node and the Base Node is lost), the Base Node will deallocate the Destination_address allocated to the
3824 Service Node SSCP session. The Base Node will use a 4-32 CL SAP (CL_432_Leave.indication) primitive to
3825 indicate the deallocation of the Destination_address to the 4-32 application running on the Base Node

3826 5.5.4 Connection establishment data format

3827 5.5.4.1 General

3828 As described in section 5.5.3, the MAC PRIME connection data is used to transfer the Device Identifier to the
3829 Base Node and the allocated Destination_address to the Service Node SSCS session. This section describes
3830 the format used for this data.

3831 5.5.4.2 Service Node to Base Node

3832 The Service Node session passes the Device Identifier to the Base Node as part of the connection
3833 establishment request. The format of this message is shown in Table 78.

3834 **Table 78 - Connection 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.

3835 5.5.4.3 Base Node to Service Node

3836 The Base Node passes the allocated Destination_address to the Service Node session as part of the
3837 connection establishment request. It also gives its own address to the Service Node. The format of this
3838 message is shown in Table 79.

3839 **Table 79 - Connection Data sent by the Base Node**

Name	Length	Description
<i>Reserved</i>	4-bits	Reserved. Should be encoded as zero in this version of the specification.
Data.DA	12-bits	Destination_address allocated to the Service Node.
<i>Reserved</i>	4-bits	Reserved. Should be encoded as zero in this version of the specification.
Data.BA	12-bits	Base_address used by the Base Node.

3840 5.5.5 Packet format

3841 The packet formats are used as defined in IEC 61334-4-32, Clause 4, LLC Protocol Data Unit Structure
3842 (LLC_PDU).

3843 5.5.6 Service Access Point

3844 5.5.6.1 Opening and closing the Convergence layer at the Service Node

3845 5.5.6.1.1 CL_432_ESTABLISH.request

3846 This primitive is passed from the application to the 4-32 Convergence layer. It is used to open a Convergence
3847 layer session and initiate the process of registering the Device Identifier with the Base Node and the Base
3848 Node allocating a Destination_address to the Service Node session.

3849 The semantics of this primitive are as follows:

3850 *CL_432_ESTABLISH.request{ DeviceIdentifier, AE }*

3851 The Device Identifier is that of the device to be registered with the Base Node.

3852 The AE parameter indicates whether the session will be authenticated and encrypted or not.

3853 If the Device Identifier is registered and the Convergence layer session is successfully opened, the primitive
3854 CL_432_ESTABLISH.confirm is used. If an error occurs the primitive CL_432_RELEASE.confirm is used.

3855 5.5.6.1.2 CL_432_ESTABLISH.confirm

3856 This primitive is passed from the 4-32 Convergence layer to the application. It is used to confirm the
3857 successful opening of the Convergence layer session and that data may now be passed over the Convergence
3858 layer.

3859 The semantics of this primitive are as follows:

3860 *CL_432_ESTABLISH.confirm{ DeviceIdentifier, Destination_address, Base_address, AE }*

3861 The Device Identifier is used to identify which CL_432_ESTABLISH.request this CL_432_ESTABLISH.confirm is
3862 for.

3863 The Destination_address is the address allocated to the Service Node 4-32 session by the Base Node.

3864 The Base_address is the address being used by the Base Node.

3865 The AE parameter indicates whether the session will be authenticated and encrypted or not.

3866 5.5.6.1.3 CL_432_RELEASE.request

3867 This primitive is passed from the application to the 4-32 Convergence layer. It is used to close the
3868 Convergence layer and release any resources it may be holding.

3869 The semantics of this primitive are as follows:

3870 *CL_432_RELEASE.request{Destination_address}*

3871 The Destination_address is the address allocated to the Service Node 4-32 session which is to be closed.

3872 The Convergence layer will use the primitive CL_432_RELEASE.confirm when the Convergence layer session
3873 has been closed.

3874 **5.5.6.1.4 CL_432_RELEASE.confirm**

3875 This primitive is passed from the 4-32 Convergence layer to the application. The primitive tells the application
3876 that the Convergence layer session has been closed. This could be because of a CL_432_RELEASE.request or
3877 because an error has occurred, forcing the closure of the Convergence layer session.

3878 The semantics of this primitive are as follows:

3879 *CL_432_RELEASE.confirm{Destination_address, result}*

3880 The Handle identifies the session which has been closed.

3881 The result parameter has the meanings defined in Table 142.

3882 **5.5.6.2 Opening and closing the Convergence layer at the Base Node**

3883 No service access point primitives are defined at the Base Node for opening or closing the Convergence layer.
3884 None are required since the 4-32 application in the Base Node does not need to pass any information to the
3885 4-32 Convergence layer in the Base Node.

3886 **5.5.6.3 Base Node indications**

3887 **5.5.6.3.1 General**

3888 The following primitives are used in the Base Node 4-32 Convergence layer to indicate events to the 4-32
3889 application in the Base Node. They indicate when a Service Node session has joined or left the network.

3890 **5.5.6.3.2 CL_432_JOIN.indicate**

3891 *CL_432_JOIN.indicate{ Device Identifier, Destination_address, AE }*

3892 The Device Identifier is that of the device connected to the Service Node that has just joined the network.

3893 The Destination_address is the address allocated to the Service Node by the Base Node.

3894 The AE parameter indicates whether messages in this session will be authenticated and encrypted or not.

3895 **5.5.6.3.3 CL_432_LEAVE.indicate**

3896 *CL_432_LEAVE.indicate{Destination_address}*

3897 The Destination_address is the address of the Service Node session that just left the network.

3898 **5.5.6.4 Data Transfer Primitives**

3899 The data transfer primitives are used as defined in IEC 61334-4-32, sections 2.2, 2.3, 2.4 and 2.11, LLC Service
3900 Specification. As stated earlier, PRIME 432 SCS make the use of IEC61334-4-32 DL_Data service (.req, .conf,
3901 .ind) for carrying out all the data involved during data transfer. Only DL_DATA service is mandatory

3902 **5.6 IPv6 Service-Specific Convergence Sublayer (IPv6 SCS)**

3903 **5.6.1 Overview**

3904 **5.6.1.1 General**

3905 The IPv6 convergence layer provides an efficient method for transferring IPv6 packets over the PRIME
3906 network.

3907 A Service Node can pass IPv6 packets to the Base Node or directly to other Service Nodes.

3908 By default, the Base Node acts as a router between the PRIME subnet and the backbone network. All the
3909 Base Nodes must have at least this connectivity capability. Any other node inside the Subnetwork can also
3910 act as a gateway. The Base Node could also act as a NAT router. However given the abundance of IPv6
3911 addresses this is not expected. How the Base Node connects to the backbone is beyond the scope of this
3912 standard.

3913 **5.6.1.2 IPv6 unicast addressing assignment**

- 3914 • IPv6 Service Nodes (and Base Nodes) shall support the standard IPv6 protocol, as described in
3915 RFC 2460.
- 3916 • IPv6 Service Nodes (and Base Nodes) shall support the standard IPv6 addressing architecture, as
3917 described in RFC 4291.
- 3918 • IPv6 Service Nodes (and Base Nodes) shall support global unicast IPv6 addresses, link-local IPv6
3919 addresses and multicast IPv6 addresses, as described in RFC 4291.
- 3920 • IPv6 Service Nodes (and Base Nodes) shall support automatic address configuration using
3921 stateless address configuration [RFC 2462]. They may also support automatic address
3922 configuration using stateful address configuration [RFC 3315] and they may support manual
3923 configuration of IPv6 addresses. The decision of which address configuration scheme to use is
3924 deployment specific.
- 3925 • Service Node shall support DHCPv6 client, when Base Nodes have to support DHCPv6 server as
3926 described in RFC 3315 for stateless address configuration

3927 **5.6.1.3 Address management in PRIME Subnetwork**

3928 Packets are routed in PRIME Subnetwork according to the node identifier NID. Node identifier is a
3929 combination of Service Node's LNID and SID (see section 4.2). The Base Node is responsible of assigning LNID

3930 to Service Nodes. During the registration process which leads to a LNID assignment to the related Service
 3931 Node, the Base Node registers the Service Node EUI-48, and the assigned LNID together with SID.

3932 At the convergence layer level, addressing is performed using the EUI-48 of the related Service Node. The
 3933 role of the convergence sublayer is to resolve the IPv6 address into EUI-48 of the Service Node. This is done
 3934 using the address resolution service set of the Base Node.

3935 **5.6.1.4 Role of the Base Node**

3936 At the convergence sublayer level, the Base Node maintains a table containing all the IPv6 unicast addresses
 3937 and the EUI-48 related to them. One of the roles of the Base Node is to perform IPv6 to EUI-48 address
 3938 resolution. Each Service Node belonging to the Subnetwork managed by the Base Node, registers its IPv6
 3939 address and EUI-48 address with the Base Node. Other Service Nodes can then query the Base Node to
 3940 resolve an IPv6 address into a EUI-48 address. This requires the establishment of a dedicated connection to
 3941 the Base Node for address resolution, which is shared by both IPv4 and IPv6 address resolution.

3942 Optionally UDP/IPv6 headers may be compressed. Currently one header compression technique is described
 3943 in the present specification that used for transmission of IPv6 packets over IEEE 802.15.4 networks, as
 3944 defined in RFC6282. This is also known as LOWPAN_IPHC1.

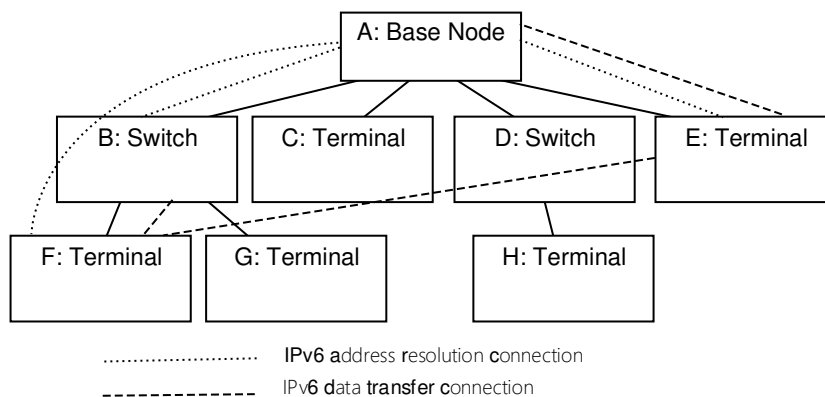
3945 The multicasting of IPv6 packets is supported using the MAC multicast mechanism

3946 **5.6.2 IPv6 Convergence layer**

3947 **5.6.2.1 Overview**

3948 **5.6.2.1.1 General**

3949 The convergence layer has a number of connection types. For address resolution there is a connection to the
 3950 Base Node. For IPv6 data transfer there are one up to many connections per destination node, depending on
 3951 how many source's and destination's IPv6 addresses are in use. This is shown in Figure 121.



3952

3953

Figure 121 - IPv6 SCS connection example

3954 Here, nodes B, E and F have address resolution connections to the Base Node. Node E has a data connection
3955 to the Base Node and node F. Node F is also has a data connection to node B. The figure does not show
3956 broadcast-traffic and multicast-traffic connections.

3957 **5.6.2.1.2 Routing in the Subnetwork**

3958 Routing IPv6 packets is the scope of the Convergence layer. In other words, the convergence layer will decide
3959 whether the packet should be sent directly to another Service Node or forwarded to the configured gateway
3960 depending on the IPv6 destination address.

3961 Although IPv6 is a connectionless protocol, the IPv6 convergence layer is connection-oriented. Once address
3962 resolution has been performed, a connection is established between the source and destination Service
3963 Nodes for the transfer of IP packets. This connection is maintained all the time the traffic is being transferred
3964 and may be removed after a period of inactivity.

3965 **5.6.2.1.3 SAR**

3966 The CPCS sublayer shall always be present with the IPv6 convergence layer allowing segmentation and
3967 reassembly facilities. The SAR sublayer functionality is given in Section 5.2. Thus, the MSDUs generated by
3968 the IPv6 convergence layer are always less than macSARSize bytes and application messages are expected to
3969 be no longer than CIMaxAppPktSize.

3970 **5.6.3 IPv6 Address Configuration**

3971 **5.6.3.1 Overview**

3972 The Service Nodes may use statically configured IPv6 addresses, link local addresses, stateless or stateful
3973 auto-configuration according to RFC 2462, or DHCPv6 to obtain IPv6 addresses. All the Nodes shall support
3974 the unicast link local address, in addition with at least one other address from those described below, and
3975 multicast addresses, if ever the node belong to multicast groups.

3976 **5.6.3.2 Interface identifier**

3977 In order to make use of stateless address auto configuration and link local addresses it is necessary to define
3978 how the Interface identifier, as defined in RFC4291, is derived. Each PRIME node has a unique EUI-48. This
3979 EUI-48 is converted into a EUI-64 in the same way as for Ethernet networks as defined in RFC2464. This EUI-
3980 64 is then used as the Interface Identifier.

3981 **5.6.3.3 IPv6 Link local address configuration**

3982 The IPv6 Link local address of a PRIME interface is formed by appending the Interface Identifier as defined
3983 above to the Prefix FE80::/64.

3984 **5.6.3.4 Stateless address configuration**

3985 An IPv6 address prefix used for stateless auto configuration, as defined in RFC4862, of a PRIME interface shall
3986 have a length of 64 bits. The IPv6 prefix is obtained by the Service Nodes from the Base Node via Router
3987 Advertisement messages (RFC 4861), which are send periodically or on request by the Base Node.

3988 **5.6.3.5 Stateful address configuration**

3989 An IPv6 address can be alternatively configured using DHCPv6, as described in RFC 3315. DHCPv6 can provide
3990 a device with addresses assigned by a DHCPv6 server and other configuration information, which are carried
3991 in options.

3992 **5.6.3.6 Multicast address**

3993 IPv6 Service Nodes (and Base Nodes) shall support the multicast IPv6 addressing, as described in RFC 4291
3994 section 2.7.

3995 **5.6.3.7 Address resolution**

3996 **5.6.3.7.1 Overview**

3997 The IPv6 layer will present the convergence layer with an IPv6 packet to be transferred. The convergence
3998 layer is responsible for determining which Service Node the packet should be delivered to, using the IPv6
3999 addresses in the packet. The convergence layer shall then establish a connection to the destination if one
4000 does not already exist so that the packet can be transferred. Two classes of IPv6 addresses can be used and
4001 the following section describes how these addresses are resolved into PRIME EUI-48 addresses. It should be
4002 noted that IPv6 does not have a broadcast address. However broadcasting is possible using multicast all
4003 nodes addresses.

4004 **5.6.3.7.2 Unicast address**

4005 **5.6.3.7.2.1 General**

4006 IPv6 unicast addresses shall be resolved into PRIME unicast EUI-48 addresses. The Base Node maintains a
4007 central database Node of IPv6 addresses and EUI-48 addresses. Address resolution functions are performed
4008 by querying this database. The Service Node shall establish a connection to the address resolution service
4009 running on the Base Node, using the TYPE value TYPE_CL_IPv6_AR. No data should be passed in the
4010 connection establishment. Using this connection, the Service Node can use two mechanisms as defined in
4011 the present specification.

4012 **5.6.3.7.2.2 Address registration and deregistration**

4013 A Service Node uses the AR_REGISTERv6_S message to register an IPv6 address and the corresponding EUI-
4014 48 address. The Base Node will acknowledge an AR_REGISTERv6_B message. The Service Node may register
4015 multiple IPv6 addresses for the same EUI-48.

4016 A Service Node uses the AR_UNREGISTERv6_S message to unregister an IPv6 address and the corresponding
4017 EUI-48 address. The Base Node will acknowledge with an AR_UNREGISTERv6_B message.

4018 When the address resolution connection between the Service Node and the Base Node is closed, the Base
4019 Node should remove all addresses associated with that connection.

4020 **5.6.3.7.2.3 Address lookup**

4021 A Service Node uses the AR_LOOKUPv6_S message to perform a lookup. The message contains the IPv6
4022 address to be resolved. The Base Node will respond with an AR_LOOKUPv6_B message that contains an error
4023 code and, if there is no error, the EUI-48 associated with the IPv6 address. If the Base Node has multiple
4024 entries in its database Node for the same IPv6 address, the possible EUI-48 returned is undefined.

4025 It should be noted that, for the link local addresses, due to the fact that the EUI-48 can be obtained from the
4026 IPv6 address, the lookup can simply return this value by extracting it from the IPv6 address.

4027 **5.6.3.7.3 Multicast address**

4028 Multicast IPv6 addresses are mapped to connection handles (ConnHandle) by the Convergence Layer (see
4029 Table 81).

4030 **5.6.3.7.4 Retransmission of address resolution packets**

4031 The connection between the Service Node and the Base Node for address resolution is not reliable. The MAC
4032 ARQ is not used. The Service Node is responsible for making retransmissions if the Base Node does not
4033 respond in one second. It is not considered an error when the Base Node receives the same registration
4034 requests multiple times or is asked to remove a registration that does not exist. These conditions can be the
4035 result of retransmissions.

4036 **5.6.4 IPv6 Packet Transfer**

4037 **5.6.4.1 Unicast transfer**

4038 **5.6.4.1.1 Gateway info**

4039 The two endpoints exchanging unicast IPv6 datagrams may both be placed within a PRIME network, or one
4040 of them can be outside. In the latter case a PRIME node shall act as gateway to route the IPv6 traffic in and
4041 out the PRIME network. It shall take 6LoWPAN encoded outgoing datagrams and reconstruct the full IPv6
4042 counterparts before sending them to the outside; and, viceversa, it shall take incoming IPv6 datagrams and
4043 6LoWPAN encoding them before injecting them into the PRIME network.

4044 The IPv6 layer registers its addresses using the primitive CL_IPv6_Register.request (5.6.9.3.2). For each one,
4045 the subnet mask and the gateway IPv6 address are provided. As soon as the address registration (5.6.4.1.2)
4046 to the BN is successful, the SSCS must store these info in an internal table like the one in Table 80.

4047

4048

Table 80 - Node addresses table entry

4049

4050

4051

4052

4053

4054

4055

Parameter	Description
CL_IPv6_Addr.Local_IP	Node's registered IPv6 address.
CL_IPv6_Addr.Subnet_Mask	Subnet mask of the IPv6 subnetwork the node IPv6 address belongs to.
CL_IPv6_Addr.Gateway_IP	IPv6 address of the subnet's gateway.

4056

5.6.4.1.2 Both endpoints in the same PRIME network

4057

4058

4059

4060

4061

4062

4063

For packets to be transferred, a connection needs to be established between the source and destination nodes. Considering that any node may be given with one to many IPv6 addresses, a different connection shall be established between two IPv6-communicating PRIME nodes for each pair of their IPv6 addresses exchanging datagrams. The IPv6 convergence layer will examine each IP packet to determine the pair of source and destination IPv6 addresses. By matching the destination IPv6 address against the subnet mask stored in Table 80 for the datagram's source IPv6 address, the SSCS determines that both endpoints are two PRIME nodes in the same PRIME network:

4064

source IPv6 addr & CL_IPv6_Addr.Subnet_Mask == destination IPv6 addr & CL_IPv6_Addr.Subnet_Mask

4065

4066

4067

If a connection between PRIME nodes owning these addresses has already been established by the SSCS, the packet is simply sent over that connection. To verify this, the convergence layer keeps a table for each connection it has with information shown in Table 81.

4068

Table 81 - IPv6 convergence layer table entry

Parameter	Description
CL_IPv6_Con.Local_IP	Local IP address of this connection
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	Header Compression scheme being used
-------------	--------------------------------------

4069 The convergence layer may close a connection when it has not been used for an implementation-defined
 4070 time period. When the connection is closed the entry for the connection is removed at both ends of the
 4071 connection.

4072 When a connection to the destination does not exist, more work is necessary. The address resolution service
 4073 is used to determine the EUI-48 address of the remote IP address. When the Base Node replies with the EUI-
 4074 48 address of the destination Service Node, a MAC connection is established to the remote device. The TYPE
 4075 value of this connection is TYPE_CL_IPv6_DATA. The data passed in the request message is defined in section
 4076 5.6.8.3. Both the local source IP address and the destination remote IPv6 address are provided so that the
 4077 remote device can properly set up its own IPv6 convergence layer table entry, thus adding the new
 4078 connection to its cache of connections for sending data in the opposite direction. Once the connection has
 4079 been established, the IP packet can be sent.

4080 **5.6.4.1.3 One endpoint outside the PRIME network**

4081 If the communication involves an endpoint external to the PRIME network, datagrams shall pass through a
 4082 border PRIME node acting as the gateway. In this scenario a PRIME connection shall be established between
 4083 the gateway and the PRIME node, and 6LoWPAN compressed datagrams shall be transmitted over it.

4084 Following subsections cover the connection establishment and usage in the two possible scenarios,
 4085 depending on which endpoint sends the first datagram.

4086 **5.6.4.1.3.1 Connection initiated by the PRIME endpoint**

4087 As for the PRIME internal data exchange (5.6.4.1.2), SSCS discovers whether the destination is outside the
 4088 PRIME network by matching the destination IPv6 address against the netmask associated with the source
 4089 IPv6 address in Table 80:

4090 *source IPv6 addr & CL_IPv6_Addr.Subnet_Mask != destination IPv6 addr & CL_IPv6_Addr.Subnet_Mask*

4091 Once the gateway IPv6 address is retrieved from the Table 80's row of the source IPv6 address, SSCS looks
 4092 up the addresses couple (source IPv6 address, gateway IPv6 address) in the Table 81 to get the connection
 4093 ID with the gateway – if a connection has not been established yet, same procedures (address resolution,
 4094 connection establishment and entry addition in Table 81) as with PRIME network internal communication
 4095 (5.6.4.1.2) shall be undergone, using the gateway IPv6 address as the remote endpoint address. Once the
 4096 connection ID is retrieved, the 6LoWPAN encoded version of the datagram shall be sent over it.

4097 **5.6.4.1.3.2 Connection initiated by the external endpoint**

4098 The SSCS discovers that the datagram is coming from the outside by failing at looking up the source IPv6
 4099 address in Table 80.

4100 By scanning rows in Table 80, the SSCS finds the gateway IPv6 address the PRIME node owning the destination
 4101 IPv6 address refers to. It is the IPv6 address stored in the table’s entry for which the following relationship
 4102 holds:

4103
$$CL_IPv6_Addr.Gateway_IP \& CL_IPv6_Addr.Subnet_Mask ==$$

 4104
$$destination\ IPv6\ addr \& CL_IPv6_Addr.Subnet_Mask$$

4105 Once the gateway IPv6 address is retrieved, SSCS looks up the addresses couple (gateway IPv6 address,
 4106 destination IPv6 address) in the Table 81 to get the connection ID with the destination PRIME node – if a
 4107 connection has not been established yet, same procedures (address resolution, connection establishment
 4108 and entry addition in Table 81) as with PRIME network internal communication (5.6.4.1.2) shall be undergone,
 4109 using the gateway IPv6 address as the local endpoint address. Once the connection ID is retrieved, the
 4110 6LoWPAN encoded version of the datagram shall be sent over it.

4111 **5.6.4.2 Multicast transfer**

4112 To join a multicast group, CL uses the MAC_JOIN.request primitive with the IPv6 address specified in the data
 4113 field. A corresponding MAC_JOIN.Confirm primitive will be generated by the MAC after completion of the
 4114 join process. The MAC_Join.Confirm primitive will contain the result (success/failure) and the corresponding
 4115 ConnHandle to be used by the CL. The MAC layer will handle the transfer of data for this connection using
 4116 the appropriate LCIDs. To leave the multicast group, the CL at the service node shall use the MAC-
 4117 LEAVE.Request{ConnHandle} primitive.

4118 To send an IPv6 multicast packet, the CL will simply send the packet to the group, using the allocated
 4119 ConnHandle.

4120 **Table 82 - IPv6 convergence layer multicast table entry**

Parameter	Description
CL_IPv6_Mul.Address	Multicast IPv6 address for this connection
CL_IPv6_Mul.ConHandle	MAC Connection handle for the connection

4121 **5.6.5 Segmentation and reassembly**

4122 The IPv6 convergence layer should support IPv6 packets with an MTU of at least 1500 bytes. This requires
 4123 the use of the common part convergence sublayer segmentation and reassembly service.

4124 **5.6.6 Compression**

4125 Any PRIME device being compliant with this Service-Specific Convergence Sublayer shall be able to decode
 4126 any valid 6LoWPAN encoded packet.

4127 All the Service Nodes and the Base Node shall support IPv6 Header Compression using source and destination
 4128 Addresses stateless compression as defined in RFC 6282. Source and destination IPv6 addresses using stateful
 4129 compression shall also be supported, but the way contexts are shared is outside the scope of this document.

4130 As far as the stateless compression of either source address or unicast destination address is concerned, the
 4131 6LoWPAN implementation in this Service-Specific Convergence Sublayer shall allow following modes only
 4132 (refer to RFC 6282 for fields' meaning):

- 4133 • Full address is carried inline (e.g. no compression) (SAM=0b00 and/or DAM=0b00).
- 4134 • Address is fully elided (SAM=0b11 and/or DAM=0b11).

4135 Remaining two modes – address compressed down to 64 bits (SAM=0b01 and/or DAM=0b01) and down to
 4136 16 bits (SAM=0b10 and/or DAM=0b10) – are forbidden as PRIME own characteristics don't allow for such
 4137 kind of compression. The full address compression (SAM=0b11 and/or DAM=0b11) is enabled by the
 4138 information added in the Table 82 once the connection between the two is established. Such a table enables
 4139 a direct mapping between the connection handle and the IPv6 addresses. When a node receives a PRIME
 4140 packet, it uses its connection handle as table's lookup key: once found, it can retrieve both IPv6 source and
 4141 destination addresses and hence restore them in the IPv6 datagram. If both IPv6 endpoints are in the same
 4142 PRIME network, both their addresses may be fully compressed by the sender. When instead one node is
 4143 outside the PRIME network, and hence the PRIME connection involves the gateway, the outer node's IPv6
 4144 address cannot be compressed, because such an address is not the one of the PRIME node being involved in
 4145 the connection; in such a situation, only the IPv6 address of the PRIME endpoint can be fully elided.

4146 **5.6.7 Quality of Service Mapping**

4147 The PRIME MAC specifies that the contention-based access mechanism supports 4 priority levels (1-4). Level
 4148 1 is used for MAC control messages, but not exclusively so.

4149 IPv6 packets include a Traffic Class field in the header to indicate the QoS the packet would like to receive.
 4150 This traffic class can be used in the same way that IPv4 TOS (see [7]). That is, three bits of the TOS indicate
 4151 the IP Precedence. The following table specifies how the IP Precedence is mapped into the PRIME MAC
 4152 priority.

4153 **Table 83 - Mapping Ipv6 precedence to PRIME MAC priority**

IP Precedence	MAC Priority
000 – Routine	3
001 – Priority	3
010 – Immediate	2

011 – Flash	2
100 – Flash Override	1
101 – Critical	1
110 – Internetwork Control	0
111 – Network Control	0

4154 **5.6.8 Packet formats and connection data**

4155 **5.6.8.1 Overview**

4156 This section defines the format of convergence layer PDUs.

4157 **5.6.8.2 Address resolution PDU**

4158 **5.6.8.2.1 General**

4159 The following PDUs are transferred over the address resolution connection between the Service Node and
 4160 the Base Node. The following sections define a number of AR.MSG values. All other values are reserved for
 4161 later versions of this standard.

4162 **5.6.8.2.2 AR_REGISTERv6_S**

4163 Table 84 shows the address resolution register message sent from the Service Node to the Base Node.

4164 **Table 84 - AR_REGISTERv6_S message format**

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type <ul style="list-style-type: none"> • For AR_REGISTERv6_S = 16
AR.IPv6	128-bits	IPv6 address to be registered
AR.EUI-48	48-bits	EUI-48 to be registered

4165 **5.6.8.2.3 AR_REGISTERv6_B**

4166 Table 85 shows the address resolution register acknowledgment message sent from the Base Node to the
4167 Service Node.

4168 **Table 85 - AR_REGISTERv6_B message format**

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type <ul style="list-style-type: none"> • For AR_REGISTERv6_B = 17
AR.IPv6	128-bits	IPv6 address registered
AR.EUI-48	48-bits	EUI-48 registered

4169 The AR.IPv6 and AR.EUI-48 fields are included in the AR_REGISTERv6_B message so that the Service Node
4170 can perform multiple overlapping registrations.

4171 **5.6.8.2.4 AR_UNREGISTERv6_S**

4172 Table 86 shows the address resolution unregister message sent from the Service Node to the Base Node.

4173 **Table 86 - AR_UNREGISTERv6_S message format**

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type <ul style="list-style-type: none"> • For AR_UNREGISTERv6_S = 18
AR.IPv6	128-bits	IPv6 address to be unregistered
AR.EUI-48	48-bits	EUI-48 to be unregistered

4174 **5.6.8.2.5 AR_UNREGISTERv6_B**

4175 Table 87 shows the address resolution unregister acknowledgment message sent from the Base Node to the
4176 Service Node.

4177

Table 87 - AR_UNREGISTERv6_B message format

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type <ul style="list-style-type: none"> For AR_UNREGISTERv6_B = 19
AR.IPv6	128-bits	IPv6 address unregistered
AR.EUI-48	48-bits	EUI-48 unregistered

4178 The AR.IPv6 and AR.EUI-48 fields are included in the AR_UNREGISTERv6_B message so that the Service Node
4179 can perform multiple overlapping unregistrations.

4180 5.6.8.2.6 AR_LOOKUPv6_S

4181 Table 888 shows the address resolution lookup message sent from the Service Node to the Base Node.

4182

Table 888 - AR_LOOKUPv6_S message format

Name	Length	Description
AR.MSG	8-bits	Address Resolution Message Type <ul style="list-style-type: none"> For AR_LOOKUPv6_S = 20
AR.IPv6	128-bits	IPv6 address to lookup

4183 5.6.8.2.7 AR_LOOKUPv6_B

4184 Table 899 shows the address resolution lookup response message sent from the Base Node to the Service
4185 Node.

4186

Table 899 - 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

Name	Length	Description
AR.IPv6	128-bits	IPv6 address looked up
AR.EUI-48	48-bits	EUI-48 for IPv6 address
AR.Status	8-bits	Lookup 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

4187 The lookup may fail if the requested address has not been registered. In that case, AR.Status will have a value
 4188 equal to 1, and the contents of AR.EUI-48 will be undefined. The lookup is only successful when AR.Status is
 4189 zero. In that case, the EUI-48 field contains the resolved address.

4190 **5.6.8.3 IPv6 Packet format**

4191 **5.6.8.3.1 General**

4192 The following PDU formats are used for transferring IPv6 packets between Service Nodes.

4193 **5.6.8.3.2 No header compression**

4194 When no header compression is used, the IP packet is simply sent as it is, without any header.

4195 **Table 90 - IPv6 Packet format without header compression**

Name	Length	Description
IPv6.PKT	n-octets	The IPv6 Packet

4196 **5.6.8.3.3 Header compression**

4197 When LOWPAN_IPHC1 header compression is used, the UDP/IPv6 packet is sent as shown in Table 91.

4198

Table 91 - 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

4199 Note that these fields are not necessarily aligned to byte boundaries. For example the IPv6.ncIPv6 field can
 4200 be any number of bits. The IPv6.IPHC_UDP field follows directly afterwards, without any padding. Padding is
 4201 only applied at the end of the complete compressed UDP/IPv6 header such that the UDP data is byte aligned.

4202 When the IPv6 packet contains data other than UDP the following packet format is used as shown in Table
 4203 92.

4204

Table 92 - 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
<i>Padding</i>	0.m-octets	Padding to byte boundary
<i>IPv6.DATA</i>	n-octets	IP Data

4205 **5.6.8.4 Connection data**

4206 **5.6.8.4.1 Overview**

4207 When a connection is established between Service Nodes for the transfer of IP packets, data is also
 4208 transferred in the connection request packets. This data allows the negotiation of compression and
 4209 notification of the IP address.

4210 **5.6.8.4.2 Unicast connection data from the initiator**

4211 Table 93 shows the connection data sent by the initiator.

4212 **Table 93 - IPv6 Unicast connection data sent by the initiator**

Name	Length	Description
Data.localIPv6	128-bits	Local IPv6 address
Data.remoteIPv6	128-bits	Remote IPv6 address

4213

4214 If the device accepts the connection, it should copy both Data.localIPv6 and Data.remoteIPv6 addresses into
 4215 a new table entry, respectively in CL_IPv6_Con.Remote_IP and CL_IPv6_Con.Local_IP.

4216 **5.6.8.4.3 Unicast connection data from the responder**

4217 No information is carried in the response's CON.DATA field.

4218 **5.6.8.4.4 Multicast connection data from the initiator**

4219 **Table 94 - IPv6 Multicast connection data sent by the initiator**

Name	Length	Description
Data.IPv6	128-bits	IPv6 multicast address

4220 It includes only IPv6 multicast address.

4221 **5.6.8.4.5 Multicast connection data from the responder**

4222 No information is carried in the response's CON.DATA field.

4223 5.6.9 Service access point

4224 5.6.9.1 Overview

4225 This section defines the service access point used by the IPv6 layer to communicate with the IPv6
4226 convergence layer.

4227 5.6.9.2 Opening and closing the convergence layer

4228 The following primitives are used to open and close the convergence layer. The convergence layer may be
4229 opened once only. The IPv6 layer may close the convergence layer when the IPv6 interface is brought down.

4230 The convergence layer will also close the convergence layer when the underlying MAC connection to the
4231 Base Node has been lost.

4232 5.6.9.2.1 CL_IPv6_Establish.request

4233 The CL_IPv6_ESTABLISH.request primitive is passed from the IPv6 layer to the IPv6 convergence layer. It is
4234 used when the IPv6 layer brings the interface up.

4235 The semantics of this primitive are as follows:

4236 *CL_IPv6_ESTABLISH.request{AE}*

4237 The AE parameter indicates whether the interface will be authenticated and encrypted or not.

4238 On receiving this primitive, the convergence layer will form the address resolution connection to the Base
4239 Node.

4240 5.6.9.2.2 CL_IPv6_Establish.confirm

4241 The CL_IPv6_ESTABLISH.confirm primitive is passed from the IPv6 convergence layer to the IPv6 layer. It is
4242 used to indicate that the convergence layer is ready to access IPv6 packets to be sent to peers.

4243 The semantics of this primitive are as follows:

4244 *CL_IPv6_ESTABLISH.confirm{AE}*

4245 The AE parameter indicates whether the interface will be authenticated and encrypted or not.}

4246 Once the convergence layer has established all the necessary connections and is ready to transmit and
4247 receive IPv6 packets, this primitive is passed to the IPv6 layer. If the convergence layer encounters an error
4248 while opening, it responds with a CL_IPv6_RELEASE.confirm primitive, rather than a
4249 CL_IPv6_ESTABLISH.confirm.

4250 **5.6.9.2.3 CL_IPv6_Release.request**

4251 The CL_IPv6_RELEASE.request primitive is used by the IPv6 layer when the interface is put down. The
4252 convergence layer closes all connections so that no more IPv6 packets are received and all resources are
4253 released.

4254 The semantics of this primitive are as follows:

4255 *CL_IPv6_RELEASE.request{}*

4256 Once the convergence layer has released all its connections and resources it returns a
4257 CL_IPv6_RELEASE.confirm.

4258 **5.6.9.2.4 CL_IPv6_Release.confirm**

4259 The CL_IPv6_RELEASE.confirm primitive is used by the IPv6 convergence layer to indicate to the IPv6 layer
4260 that the convergence layer has been closed. This can be as a result of a CL_IPv6_RELEASE.request primitive,
4261 a CL_IPv6_ESTABLISH.request primitive, or because the MAC layer indicates the address resolution
4262 connection has been lost, or the Service Node itself is no longer registered.

4263 The semantics of this primitive are as follows:

4264 *CL_IPv6_RELEASE.confirm{result}*

4265 The result parameter has the meanings defined in Table 142.

4266 **5.6.9.3 Unicast address management**

4267 **5.6.9.3.1 General**

4268 The primitives defined here are used for address management, i.e. the registration and unregistration of IPv6
4269 addresses associated with this convergence layer.

4270 When there are no IPv6 addresses associated with the convergence layer, the convergence layer will only
4271 send and receive multicast packets; unicast packets may not be sent. However, this is sufficient for various
4272 address discovery protocols to be used to gain an IPv6 address. Once an IPv6 address has been registered,
4273 the IPv6 layer can transmit unicast packets that have a source address equal to one of its registered
4274 addresses.

4275 **5.6.9.3.2 CL_IPv6_Register.request**

4276 This primitive is passed from the IPv6 layer to the IPv6 convergence layer to register an IPv6 address.

4277 The semantics of this primitive are as follows:

4278 *CL_IPv6_REGISTER.request{ipv6, netmask, gateway}*

4279 The ipv6 address is the address to be registered.

4280 The netmask is the network mask, used to mask the network number from the address. The netmask is used
4281 by the convergence layer to determine whether the packet should deliver directly or the gateway should be
4282 used.

4283 The IPv6 address of the gateway, to which packets with destination address that are not in the same subnet
4284 as the local address are to be sent.

4285 Once the IPv6 address has been registered to the Base Node, a CL_IPv6_REGISTER.confirm primitive is used.
4286 If the registration fails, the CL_IPv6_RELEASE.confirm primitive will be used.

4287 **5.6.9.3.3 CL_IPv6_Register.confirm**

4288 This primitive is passed from the IPv6 convergence layer to the IPv6 layer to indicate that a registration has
4289 been successful.

4290 The semantics of this primitive are as follows:

4291 *CL_IPv6_REGISTER.confirm{ipv6}*

4292 The ipv6 address is the address that was registered.

4293 Once registration has been completed, the IPv6 layer may send IPv6 packets using this source address.

4294 **5.6.9.3.4 CL_IPv6_Unregister.request**

4295 This primitive is passed from the IPv6 layer to the IPv6 convergence layer to unregister an IPv6 address.

4296 The semantics of this primitive are as follows:

4297 *CL_IPv6_UNREGISTER.request{ipv6}*

4298 The ipv6 address is the address to be unregistered.

4299 Once the IPv6 address has been unregistered to the Base Node, a CL_IPv6_UNREGISTER.confirm primitive is
4300 used. If the registration fails, the CL_IPv6_RELEASE.confirm primitive will be used.

4301 **5.6.9.3.5 CL_IPv6_Unregister.confirm**

4302 This primitive is passed from the IPv6 convergence layer to the IPv6 layer to indicate that an unregistration
4303 has been successful.

4304 The semantics of this primitive are as follows:

4305 *CL_IPv6_UNREGISTER.confirm{ipv6}*

4306 The IPv6 address is the address that was unregistered.

4307 Once unregistration has been completed, the IPv6 layer may not send IPv6 packets using this source address.

4308 **5.6.9.4 Multicast group management**

4309 **5.6.9.4.1 General**

4310 This section describes the primitives used to manage multicast groups.

4311 **5.6.9.4.2 CL_IPv6_MUL_Join.request**

4312 This primitive is passed from the IPv6 layer to the IPv6 convergence layer. It contains an IPv6 multicast
4313 address that is to be joined.

4314 The semantics of this primitive are as follows:

4315 *CL_IPv6_MUL_JOIN.request{IPv6, AE}*

4316 The IPv6 address is the IPv6 multicast group that is to be joined.

4317 The AE parameter indicates whether messages in this group will be authenticated and encrypted or not.

4318 When the convergence layer receives this primitive, it will arrange for IP packets sent to this group to be
4319 multicast in the PRIME network and receive packets using this address to be passed to the IPv6 stack. If the
4320 convergence layer cannot join the group, it uses the CL_IPv6_MUL_LEAVE.confirm primitive. Otherwise the
4321 CL_IPv6_MUL_JOIN.confirm primitive is used to indicate success.

4322 **5.6.9.4.3 CL_IPv6_MUL_Join.confirm**

4323 This primitive is passed from the IPv6 convergence layer to the IPv6. It contains a result status and an IPv6
4324 multicast address that was joined.

4325 The semantics of this primitive are as follows:

4326 *CL_IPv6_MUL_JOIN.confirm{IPv6, AE}*

4327 The IPv6 address is the IPv6 multicast group that was joined. The convergence layer will start forwarding IPv6
4328 multicast packets for the given multicast group.

4329 The AE parameter indicates whether messages in this group will be authenticated and encrypted or not.

4330 **5.6.9.4.4 CL_IPv6_MUL_Leave.request**

4331 This primitive is passed from the IPv6 layer to the IPv6 convergence layer. It contains an IPv6 multicast
4332 address to be left.

4333 The semantics of this primitive are as follows:

4334 *CL_IPv6_MUL_LEAVE.request{IPv6}*

4335 The IPv6 address is the IPv6 multicast group to be left. The convergence layer will stop forwarding IPv6
4336 multicast packets for this group and may leave the PRIME MAC multicast group.

4337 **5.6.9.4.5 CL_IPv6_MUL_Leave.confirm**

4338 This primitive is passed from the IPv6 convergence layer to the IPv6. It contains a result status and an IPv6
4339 multicast address that was left.

4340 The semantics of this primitive are as follows:

4341 *CL_IPv6_MUL_LEAVE.confirm{IPv6, Result}*

4342 The IPv6 address is the IPv6 multicast group that was left. The convergence layer will stop forwarding IPv6
4343 multicast packets for the given multicast group.

4344 The Result takes a value from Table 142.

4345 This primitive can be used by the convergence layer as a result of a CL_IPv6_MUL_JOIN.request,
4346 CL_IPv6_MUL_LEAVE.request or because of an error condition resulting in the loss of the PRIME MAC
4347 multicast connection.

4348 **5.6.9.5 Data transfer**

4349 **5.6.9.5.1 General**

4350 The following primitives are used to send and receive IPv6 packets.

4351 **5.6.9.5.2 CL_IPv6_DATA.request**

4352 This primitive is passed from the IPv6 layer to the IPv6 convergence layer. It contains one IPv6 packet to be
4353 sent.

4354 The semantics of this primitive are as follows:

4355 *CL_IPv6_DATA.request{IPv6_PDU}*

4356 The IPv6_PDU is the IPv6 packet to be sent.

4357 **5.6.9.5.3 CL_IPv6_DATA.confirm**

4358 This primitive is passed from the IPv6 convergence layer to the IPv6 layer. It contains a status indication and
4359 an IPv6 packet that has just been sent.

4360 The semantics of this primitive are as follows:

4361 *CL_IPv6_DATA.confirm{IPv6_PDU, Result}*

4362 The IPv6_PDU is the IPv6 packet that was to be sent.

4363 The Result value indicates whether the packet was sent or an error occurred. It takes a value from Table 142.

4364 **5.6.9.5.4 CL_IPv6_DATA.indicate**

4365 This primitive is passed from the IPv6 convergence layer to the IPv6 layer. It contains an IPv6 packet that has
4366 just been received.

4367 The semantics of this primitive are as follows:

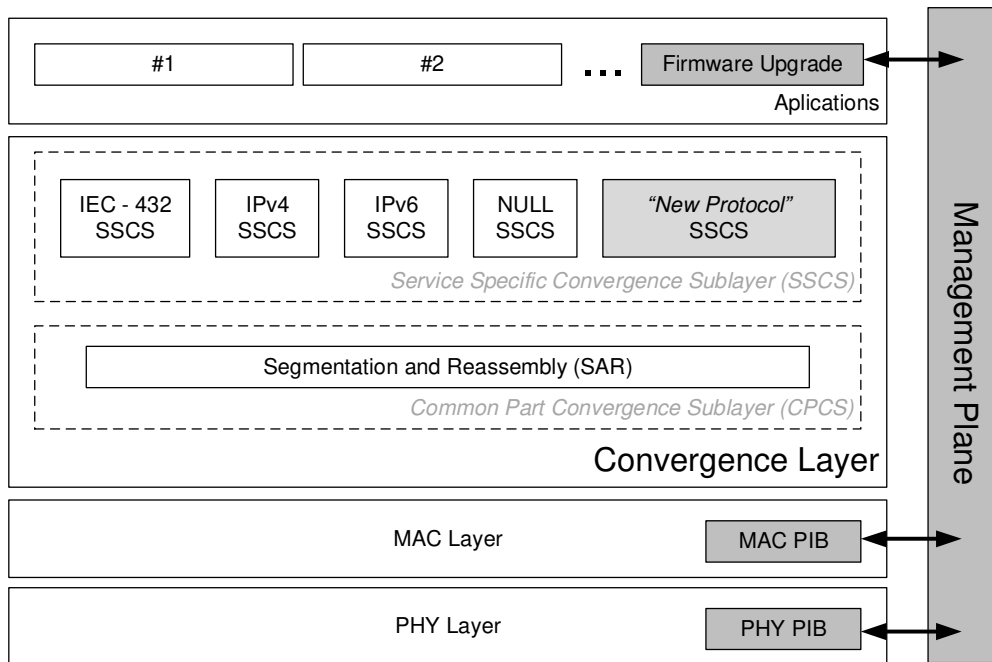
4368 *CL_IPv6_DATA.indicate{IPv6_PDU }*

4369 The IPv6_PDU is the IPv6 packet that was received.

4370 6 Management plane

4371 6.1 Introduction

4372 This chapter specifies the Management plane functionality. The picture below highlights the position of
4373 Management plane in overall protocol architecture.



4374

4375

Figure 122 - Management plane. Introduction.

4376 All nodes shall implement the management plane functionality enumerated in this section. Management
4377 plane enables a local or remote control entity to perform actions on a Node.

4378 Present version of this specification enumerates management plane functions for Node management and
4379 firmware upgrade. Future versions may include additional management functions.

- 4380 • To enable access to management functions on a Service Node, Base Node shall open a
- 4381 management connection after successful completion of registration (refer to 6.4)
- 4382 • The Base Node may open such a connection either immediately on successful registration or
- 4383 sometime later.
- 4384 • Unicast management connection shall be identified with CON.TYPE = TYPE_CL_MGMT.
- 4385 • Multicast management connections can also exist. At the time of writing of this document,
- 4386 multicast management connection shall only be used for firmware upgrade.
- 4387 • There shall be no broadcast management connection.
- 4388 • In case Service Node supports ARQ connections, the Base Node shall preferentially try to open
- 4389 an ARQ connection for management functions.
- 4390 • Management plane functions shall use NULL SSCS as specified in section 0

4391 6.2 Node management

4392 6.2.1 General

4393 Node management is accomplished through a set of attributes. Attributes are defined for both PHY and MAC
4394 layers. The set of these management attributes is called PLC Information Base (PIB). Some attributes are
4395 read-only while others are read-write.

4396 PIB Attribute identifiers are 16 bit values. This allows for up to 65535 PIB Attributes to be specified.

- 4397 • PIB Attribute identifier values from 0 to 32767 are open to be standardized. No proprietary
4398 attributes may have identifiers in this range.
- 4399 • Values in the range 32768 to 65535 are open for vendor specific usage.

4400 PIB Attributes identifiers in standard range (0 to 32767) that are not specified in this version are reserved for
4401 future use.

4402 **Note:** PIB attribute tables below indicate type of each attribute. For integer types the size of the integer has
4403 been specified in bits. An implementation may use a larger integer for an attribute; however, it must not use
4404 a smaller size.

4405 6.2.2 PHY PIB attributes

4406 6.2.2.1 General

4407 The PHY layer implementation in each device may optionally maintain a set of attributes which provide
4408 detailed information about its working. The PHY layer attributes are part of the PLC Information Base (PIB).

4409 6.2.2.2 Statistical attributes

4410 The PHY may provide statistical information for management purposes. Next table lists the statistics that PHY
4411 should make available to management entities across the PLME_GET primitive. The Id field in this table is the
4412 service parameter of the PLME_GET primitive specified in section 3.11.4.

4413 **Table 95 - 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

Attribute Name	Size (in bits)	Id	Description
			PHY header had an invalid value. This count would reflect number of times corrupt data was received and the CRC calculation failed to detect it.
phyStatsTxDropCount	16	0x00A2	Number of times when PHY layer received new data to transmit (PHY_DATA.request) and had to either overwrite on existing data in its transmit queue or drop the data in new request due to full queue.
phyStatsRxDropCount	16	0x00A3	Number of times when PHY layer received new data on the channel and had to either overwrite on existing data in its receive queue or drop the newly received data due to full queue.
phyStatsRxTotalCount	32	0x00A4	Total number of PPDU's correctly decoded. Useful for PHY layer test cases, to estimate the FER.
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 exponential moving average (EMA). Next value is $V_{next} = S * NewSample + (1-S) * V_{prev}$ Where $S = 1 / (2^{phyEMASmoothing}) .$

4414 **6.2.2.3 Implementation attributes**

4415 It is possible to implement PHY functions conforming to this specification in multiple ways. The multiple
 4416 implementation options provide some degree of unpredictability for MAC layers. PHY implementations may
 4417 optionally provide specific information on parameters which are of interest to MAC across the PLME_GET
 4418 primitive. A list of such parameters which maybe queried across the PLME_GET primitives by MAC is provided
 4419 in Table 96. All of the attributes listed in Table 96 are implementation constants and shall not be changed.

4420

Table 96 - PHY read-only parameters, providing information on specific implementation

Attribute Name	Size (in bits)	Id	Description
phyTxQueueLen	10	0x00B0	Number of concurrent MPDUs that the PHY transmit buffers can hold.
phyRxQueueLen	10	0x00B1	Number of concurrent MPDUs that the PHY receive buffers can hold.
phyTxProcessingDelay	20	0x00B2	Time elapsed from the instance when data is received on 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 physical channel to the time when it is made available to 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.
phyAgcStepValue	3	0x00B5	Distance between steps in dB \leq 6dB.
phyAgcStepNumber	8	0x00B6	Number of steps so that $\text{phyAgcMinGain} + (\text{phyAgcStepNumber} - 1) * \text{phyAgcStepValue} \geq 21\text{dB}$.

4421 **6.2.3 MAC PIB attributes**4422 **6.2.3.1 General**

4423 **Note:** Note that the “M”(Mandatory) column in the tables below specifies if the PIB attributes are mandatory
 4424 for all devices (both Service Node and Base Node, specified as “All”), only for Service Nodes (“SN”), only for
 4425 Base Nodes (“BN”) or not mandatory at all (“No”).

4426 **6.2.3.2 MAC variable attributes**

4427 MAC PIB variables include the set of PIB attributes that influence the functional behavior of an
 4428 implementation. These attributes may be defined external to the MAC, typically by the management entity
 4429 and implementations may allow changes to their values during normal running, i.e. even after the device
 4430 start-up sequence has been executed.

4431 An external management entity can have access to these attributes through the MLME_GET (4.5.5.7) and
 4432 MLME_SET (4.5.5.9) set of primitives. The Id field in the following table would be the *PIBAttribute* that needs
 4433 to be passed MLME SAP while working on these parameters

4434

Table 97 - Table of MAC read-write variables

Attribute Name	Id	Type	M	Valid Range	Description	Def.
macVersion	0x0001	Integer8	All	0x01	The current MAC Version. This is a ‘read-only’ attribute	0x01
macMinSwitchSearchTime	0x0010	Integer8	No	16 – 32 seconds	Minimum time for which a Service Node in <i>Disconnected</i> status should scan the channel for Beacons before it can broadcast PNPDU. This attribute is not maintained in Base Nodes.	24

Attribute Name	Id	Type	M	Valid Range	Description	Def.
macMaxPromotionPdu	0x0011	Integer8	No	1 – 4	<p>Maximum number of PNPDU's that may be transmitted by a Service Node in a period of <i>macPromotionPduTxPeriod</i> seconds.</p> <p>This attribute is not maintained in Base Node.</p>	2
macPromotionPduTxPeriod	0x0012	Integer8	No	2 – 8 seconds	<p>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.</p>	5
macSCPMaxTxAttempts	0x0014	Integer8	No	2 – 5	<p>Number of times the CSMA algorithm would attempt to transmit requested data when a previous attempt was withheld due to PHY indicating channel busy.</p>	5
macMinCtlReTxTimer	0x0015	Integer8	All	2 sec	<p>Minimum number of seconds for which a MAC entity waits for acknowledgement of receipt of MAC Control Packet from its peer entity. On expiry of this time, the MAC entity may retransmit the MAC Control Packet.</p>	2

Attribute Name	Id	Type	M	Valid Range	Description	Def.
macCtrlMsgFailTime	0x0018	Integer8	No	6 - 100	Number of seconds for which a MAC entity in Switch Nodes waits before declaring a children's transaction procedures expired	45
macEMASmoothing	0x0019	Integer8	All	0 - 7	Smoothing factor divider for values that are updated as exponential moving average (EMA). Next value is $V_{next} = S * NewSample + (1 - S) * V_{prev}$ Where $S = 1 / (2^{macEMASmoothing})$	3
macMinBandSearchTime	0x001A	Integer8	No	32 – 120	Period of time in seconds for which a disconnected Node listen on a specific band before moving to one other.	60
macPromotionMaxTxPeriod	0x001B	Integer8	SN	16-120	Period of time in seconds for which at least one PNPDU shall be sent Note: This attribute is deprecated in v1.4 and only maintained by devices implementing v1.3.6	32
macPromotionMinTxPeriod	0x001C	Integer8	SN	2-16	Period of time in seconds for which at no more than one PNPDU shall be sent.	2

Attribute Name	Id	Type	M	Valid Range	Description	Def.
macSARSize	0x001D	Integer8	All	0-7	<p>Maximum Data packet size that can be accepted with the MCPS-DATA.Request</p> <p>0: Not mandated by BN (SAR operates normally)</p> <p>1: SAR = 16 bytes</p> <p>2: SAR =32 bytes</p> <p>3: SAR = 48 bytes</p> <p>4: SAR =64 bytes</p> <p>5: SAR =128 bytes</p> <p>6: SAR =192 bytes</p> <p>7: SAR =255 bytes</p> <p>This attribute can be modified only in Base Node. Read-only for Service Nodes</p>	0
macRobustnessManagement	0x004A	Integer8	No	0-3	<p>Force the network to operate only with one specific modulation</p> <p>0 – No forcing automatic robustness-management</p> <p>1 – Use only DBPSK_CC</p> <p>2 - Use only DQPSK_R</p> <p>3 - Use only DBPSK_R</p> <p>This attribute can be modified only in Base Node. Read-only for Service Nodes</p>	0

Attribute Name	Id	Type	M	Valid Range	Description	Def.
macUpdatedRMTimeout	0x004B	Integer16	All	60-3600	Period of time in seconds for which an entry in the	240
macALVHopRepetitions	0x004C	Integer8	All	0-7	Number of repetition for the ALV packets	5

4435

Table 98 - Table of MAC read-only variables

Attribute Name	Id	Type	M	Valid Range	Description	Def.
macSCPChSenseCount	0x0017	Integer8	No	2 – 5	Number of times for which an implementation has to perform channel-sensing. This is a 'read-only' attribute.	-
macEUI-48	0x001F	EUI-48	All		EUI-48 of the Node	-
macCSMAR1	0x0034	Integer8	All	0 - 4	Control how fast the CSMA contention window shall increase. Controls exponential increase of initial CSMA contention window size	3
macCSMAR2	0x0035	Integer8	All	1 - 4	Control initial CSMA contention window size. Controls linear increase of initial CSMA contention window size.	1
macCSMADelay	0x0038	Integer8	All	3ms – 9ms	The delay between two consecutive CSMA channel senses.	3 ms

Attribute Name	Id	Type	M	Valid Range	Description	Def.
macCSMAR1Robust	0x003B	Integer8	All	0 - 5	Control how fast the CSMA contention window shall increase when node supports Robust Mode. Controls exponential increase of initial CSMA contention window size.	4
macCSMAR2Robust	0x003C	Integer8	All	1 - 8	Control initial CSMA contention window size when node supports Robust Mode. Controls linear increase of initial CSMA contention window size.	2
macCSMADelayRobust	0x003D	Integer8	All	3ms – 9ms	The delay between two consecutive CSMA channel senses when node supports Robust Mode.	6 ms
macAliveTimeMode	0x003E	Integer8	BN	0 - 1	Selects the <i>MACAliveTime</i> value mapping for network Alive time. 0 – 1.4 mode 1 – BC mode	-

4436 6.2.3.3 Functional attributes

4437 Some PIB attributes belong to the functional behavior of MAC. They provide information on specific aspects.
4438 A management entity can only read their present value using the MLME_GET primitives. The value of these
4439 attributes cannot be changed by a management entity through the MLME_SET primitives.

4440 The Id field in the table below would be the *PIBAttribute* that needs to be passed MLME_GET SAP for
4441 accessing the value of these attributes.

4442

Table 99 - 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. (0x0000 is reserved for Base Node)
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 <i>Terminal</i> functional state. (0x00 is reserved for Base Node)
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 present frame.

Attribute Name	Id	Type	M	Valid Range	Description
macNodeHierarchyLevel	0x0026	Integer8	SN	0 – 63	Level of this Node in Subnetwork hierarchy.
macBeaconRxPos	0x0039	Integer16	SN	0 – 1104	Beacon Position on which this device's Switch Node transmits its beacon. Position is expressed in terms of symbols from the start of the frame. This attribute is not maintained in a Base Node.
macBeaconTxPos	0x003A	Integer8	SN	0 – 1104	Beacon Position in which this device transmits its beacon. Position is expressed in terms of symbols from the start of the frame. This attribute is not maintained in Service Nodes that are in a <i>Terminal</i> functional state.
macBeaconRxFrequency	0x002A	Integer8	SN	0 – 5	Number of frames between receptions of two successive beacons. A value of 0x0 indicates beacons are received in every frame. This attribute is not maintained in Base Node. Use the same encoding of FRQ field in the packets
macBeaconTxFrequency	0x002B	Integer8	SN	0 – 5	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 <i>Terminal</i> functional state. Use the same encoding of FRQ field in the packets

Attribute Name	Id	Type	M	Valid Range	Description
macCapabilities	0x002C	Integer16	All	Bitmap	<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> <p>Bit0: Robust mode Capable;</p> <p>Bit1: Backward Compatible Capable;</p> <p>Bit2: Switch Capable;</p> <p>Bit3: Packet Aggregation Capable;</p> <p>Bit4: Connection Free Period Capable;</p> <p>Bit5: Direct Connection Capable;</p> <p>Bit6: ARQ Capable;</p> <p>Bit7: Reserved for future use;</p> <p>Bit8: Direct Connection Switching;</p> <p>Bit9: Multicast Switching Capability;</p> <p>Bit10: Robust promotion device Capable;</p> <p>Bit11: ARQ Buffering Switching Capability;</p> <p>Bits12 to 15: Reserved for future use.</p>

Attribute Name	Id	Type	M	Valid Range	Description
macFrameLength	0x002D	Integer16	All	0 – 3	Frame Length in the present super-frame 0 - 276 symbols 1 - 552 symbols 2 - 828 symbols 3 - 1104 symbols
macCFPLength	0x002E	Integer16	All		The CFP length in symbols, in present frame
macGuardTime	0x002F	Integer16	All	1 symbol	The guard time between portion of the frame in symbols
macBCMode	0x0030	Integer16	All	0 or 1	MAC is operating in Backward Compatibility Mode
macBeaconRxQty	0x0032	Integer16	All		The QLTy field this device's Switch Node transmits its beacon.
macBeaconTxQty	0x0033	Integer16	All		The QLTy field this device transmits its beacon.

4443 6.2.3.4 Statistical attributes

4444 The MAC layer shall provide statistical information for management purposes. Table 100 lists the statistics
4445 MAC shall make available to management entities across the MLME_GET primitive.

4446 The Id field in table below would be the *PIBAttribute* that needs to be passed MLME_GET SAP for accessing
4447 the value of these attributes.

4448

Table 100 - 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 channel busy state.

4449

6.2.3.5 MAC list attributes

4450

MAC layer shall make certain lists available to the management entity across the MLME_LIST_GET primitive.

4451

These lists are given in Table 101. Although a management entity can read each of these lists, it cannot

4452

change the contents of any of them.

4453

The Id field in table below would be the *PIBListAttribute* that needs to be passed MLME_LIST_GET primitive

4454

for accessing the value of these attributes.

4455

Table 101 - Table of read-only lists made available by MAC layer through 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.

List Attribute Name	Id	M	Description		
			regEntryLNID	Integer16	LNID allocated to this Node.
			regEntryState	TERMINAL=1 , SWITCH=2	Functional state of this Node.
			regEntryLSID	Integer8	SID allocated to this Node.
			regEntrySID	Integer8	SID of Switch through which this Node is connected.
			regEntryLevel	Integer8	Hierarchy level of this Node.
			regEntryTCap	Integer8	<p>Bitmap of MAC Capabilities of Terminal functions in this device.</p> <p>Bits in sequence of right-to-left shall have the following meaning:</p> <p>Bit0: Robust mode Capable; Bit1: Backward Compatible Capable; Bit2: Switch Capable; Bit3: Packet Aggregation Capable; Bit4: Connection Free Period Capable; Bit5: Direct Connection Capable; Bit6: ARQ Capable; Bit7: Reserved for future use.</p>

List Attribute Name	Id	M	Description		
			regEntrySwCap	Integer8	<p>Bitmap of MAC Switching capabilities of this device</p> <p>Bits in sequence of right-to-left shall have the following meaning:</p> <p>Bit0: Direct Connection Switching;</p> <p>Bit1:Reserved;</p> <p>Bit2:Reserved;</p> <p>Bit3: ARQ Buffering Switching Capability;</p> <p>Bit4 to 7:Reserved for future use.</p>
maListActiveConn	0x0051	BN	List of active non-direct connections. This list is maintained by the Base Node only.		
			Entry Element	Type	Description
			connEntrySID	Integer8	SID of Switch through which the Service Node is connected.
			connEntryLNID	Integer16	NID allocated to Service Node.
			connEntryLCID	Integer16	LCID allocated to this connection.
			connEntryID	EUI-48	EUI-48 of Service Node.
maListMcastEntries	0x0052	No	List of entries in multicast switching table. This list is not maintained by Service Nodes in a <i>Terminal</i> functional state.		
			Entry Element	Type	Description
			mcastEntryLCID	Integer16	LCID of the multicast group.

List Attribute Name	Id	M	Description		
			mcastEntryMembers	Integer16	Number of child Nodes (including the Node itself) that are members of this group.
macListSwitchTable	0x005A	SN	List the Switch table. This list is not maintained by Service Nodes in a <i>Terminal</i> functional state.		
			Entry Element	Type	Description
			stblEntryLNID	Integer 16	LNID of attached Switch Node.
			stblEntryLSID	Integer8	LSID assigned to the attached Switch Node.
			stblEntrySID	Integer8	SID of attached Switch Node
			stblEntryALVTime	Integer8	The TIME value used for the Keep Alive process
macListDirectConn	0x0054	No	List of direct connections that are active. This list is maintained only in the Base Node.		
			Entry Element	Type	Description
			dconnEntrySrcSID	Integer8	SID of Switch through which the source Service Node is connected.
			dconEntrySrcLNID	Integer16	NID allocated to the source Service Node.
			dconnEntrySrcLCID	Integer16	LCID allocated to this connection at the source.
			dconnEntrySrcID	EUI-48	EUI-48 of source Service Node.

List Attribute Name	Id	M	Description		
			dconnEntryDstSID	Integer8	SID of Switch through which the destination Service Node is connected.
			dconnEntryDstLNID	Integer16	NID allocated to the destination Service Node.
			dconnEntryDstLCID	Integer16	LCID allocated to this connection at the destination.
			dconnEntryDstEUI48	EUI-48	EUI-48 of destination Service Node.
			dconnEntryDSID	Integer8	SID of Switch that is the direct Switch.
			dconnEntryDEUI48	EUI-48	EUI-48 of direct switch.
maListDirectTable	0x0055	No	List the direct Switch table		
			Entry Element	Type	Description
			dconnEntrySrcSID	Integer8	SID of Switch through which the source Service Node is connected.
			dconEntrySrcLNID	Integer16	NID allocated to the source Service Node.
			dconnEntrySrcLCID	Integer16	LCID allocated to this connection at the source.
			dconnEntryDstSID	Integer8	SID of Switch through which the destination Service Node is connected.
			dconnEntryDstLNID	Integer16	NID allocated to the destination Service Node.
			dconnEntryDstLCID	Integer16	LCID allocated to this connection at the destination.

List Attribute Name	Id	M	Description		
			dconnEntryDID	EUI-48	EUI-48 of direct switch.
macListAvailableSwitches	0x0056	SN	List of Switch Nodes whose beacons are received.		
			Entry Element	Type	Description
			slistEntrySNA	EUI-48	EUI-48 of the Subnetwork.
			slistEntryLSID	Integer8	SID of this Switch.
			slistEntryLevel	Integer8	Level of this Switch in Subnetwork hierarchy.
			slistEntryRxLvl	Integer8 EMA	Received signal level for this Switch.
			slistEntryRxSNR	Integer8 EMA	Signal to Noise Ratio for this Switch.
	0x0057		Deprecated since v1.3.6 of specs and reserved for future use		
macListActiveConnections	0x0058	All	List of active non-direct connections. This list is maintained by the Base Node only. Extended version.		
			Entry Element	Type	Description
			connEntrySID	Integer16	SID of Switch through which the Service Node is connected.
			connEntryLNID	Integer16	NID allocated to Service Node.
			connEntryLCID	Integer16	LCID allocated to this connection.

List Attribute Name	Id	M	Description		
			connEntryID	EUI-48	EUI-48 of Service Node.
			connType	Integer8	Type of connection.
macListPhyComm	0x0059	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 is contains also entries for every directly connected child Node.		
			Entry Element	Type	Description
			phyCommLNID	Integer16	LNID of the peer device
			phyCommSID	Integer8	SID of the peer device
			phyCommTxPwr	Integer8	Tx power of GPDU packets send to the device.
			phyCommRxLvl	Integer8 EMA	Rx power level of GPDU packets received from the device.
			phyCommSNR	Integer8 EMA	SNR of GPDU packets received from the device.
			phyCommTxModulation	Integer8	Modulation scheme to be used for communicating with this node.
			phyCommPhyTypeCapability	Integer8	Capability of the node to receive only PHY Type A or PHY Type A+B frames 0: Type A only node 1: Type A+B capable node

List Attribute Name	Id	M	Description		
			phyCommRxAge	Integer16	Time [seconds] since last update of phyCommTxModulation.

4456 **6.2.3.6 MAC security attributes**

4457 **Table 102 - MAC security attributes**

Attribute Name	Id	Size	Description
macSecDUK	0x005B	128 bits	<p>Device Unique Key to use in initial key derivation functions. The key shall be updated immediately; it shall not require re-registering the node.</p> <p>As a guideline, the Base Node should store both the old and new keys until the node has complete a successful registration with the new one, in the reception of a REG_REQ the Base Node should authenticate with the new key and if failed try again with the old one.</p> <p>Access to this PIB shall have the following restrictions:</p> <ul style="list-style-type: none"> • Is write only, shall not be read. • Shall only be available if the underlying connection is encrypted, authenticated and unicast.
MACUpdateKeysTime	0x005C	32 bits	<p>Maximum time in seconds allowed using the same SWK or WK. After that time the keys shall no longer be considered valid.</p> <p>The maximum allowed value for this PIB is defined in Table 13, the value depends on the number of channels.</p>

4458 **6.2.3.7 Action PIB attributes**

4459 Some of the conformance tests require triggering certain actions on Service Nodes and Base Nodes. The
 4460 following table lists the set of action attributes that need to be supported by all implementations.

4461

Table 103 - Action PIB attributes

Attribute Name	Id	M	Size (bytes)	Description												
MACActionTxData	0x0060	SN	1	Total number of PPDUs correctly decoded. Useful for PHY layer to estimate FER.												
MACActionConnClose	0x0061	SN	1	Trigger to close one of the open connections.												
MACActionRegReject	0x0062	SN	1	Trigger to reject incoming registration request.												
MACActionProReject	0x0063	SN	1	Trigger to reject incoming promotion request.												
MACActionUnregister	0x0064	SN	1	Trigger to unregister from the Subnetwork.												
MACActionPromote	0x0065	BN	6	Trigger to promote a given Service Node from the Subnetwork. PARAM: EUI-48 of the node being promoted.												
MACActionDemote	0x0066	BN	6	Trigger to demote a given Service Node from the Subnetwork. PARAM: EUI-48 of the node being demoted.												
MACActionReject	0x0067	BN		Rejects or stops (toggles) rejecting packets of a certain type												
				<table border="1"> <thead> <tr> <th>Entry Element</th> <th>Size</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Node</td> <td>6</td> <td>EUI 48 of the Node</td> </tr> <tr> <td>Reject</td> <td>1</td> <td>1 – Reject 0 – Stop Rejecting</td> </tr> <tr> <td>Type</td> <td>1</td> <td>0 – rejects PRO_REQ_S 1 – rejects PRM</td> </tr> </tbody> </table>	Entry Element	Size	Description	Node	6	EUI 48 of the Node	Reject	1	1 – Reject 0 – Stop Rejecting	Type	1	0 – rejects PRO_REQ_S 1 – rejects PRM
Entry Element	Size	Description														
Node	6	EUI 48 of the Node														
Reject	1	1 – Reject 0 – Stop Rejecting														
Type	1	0 – rejects PRO_REQ_S 1 – rejects PRM														

Attribute Name	Id	M	Size (bytes)	Description																																
				2 – rejects CON_REQ_S 3- rejects REG_REQ																																
MACAliveTime	0x0068	BN	1	Forces alive time for the network or sets it as automatic. <table border="1" data-bbox="799 714 1449 1514"> <thead> <tr> <th rowspan="2">Value</th> <th colspan="2">macAliveTimeMode</th> </tr> <tr> <th>0</th> <th>1</th> </tr> </thead> <tbody> <tr> <td>0x00</td> <td>128s</td> <td>32s</td> </tr> <tr> <td>0x01</td> <td>256s</td> <td>64s</td> </tr> <tr> <td>0x02</td> <td>512s</td> <td>128s</td> </tr> <tr> <td>0x03</td> <td>2048s</td> <td>256s</td> </tr> <tr> <td>0x04</td> <td>4096s</td> <td>512s</td> </tr> <tr> <td>0x05</td> <td>8192s</td> <td>1024s</td> </tr> <tr> <td>0x06</td> <td>16384s</td> <td>2048s</td> </tr> <tr> <td>0x07</td> <td>32768s</td> <td>4096s</td> </tr> <tr> <td>0xFF</td> <td colspan="2">Reset alive time to the configured value.</td> </tr> </tbody> </table>	Value	macAliveTimeMode		0	1	0x00	128s	32s	0x01	256s	64s	0x02	512s	128s	0x03	2048s	256s	0x04	4096s	512s	0x05	8192s	1024s	0x06	16384s	2048s	0x07	32768s	4096s	0xFF	Reset alive time to the configured value.	
Value	macAliveTimeMode																																			
	0	1																																		
0x00	128s	32s																																		
0x01	256s	64s																																		
0x02	512s	128s																																		
0x03	2048s	256s																																		
0x04	4096s	512s																																		
0x05	8192s	1024s																																		
0x06	16384s	2048s																																		
0x07	32768s	4096s																																		
0xFF	Reset alive time to the configured value.																																			
	0x0069			Deprecated since v1.3.6 and reserved for future use																																
MACActionBroadcastDataBurst	0x006A	BN		Send a burst of data PDU-s with a test sequence using broadcast																																
				<table border="1" data-bbox="786 1749 1465 1977"> <thead> <tr> <th>Entry Element</th> <th>Size</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Number</td> <td>4</td> <td>Number of PDUs to be sent</td> </tr> </tbody> </table>	Entry Element	Size	Description	Number	4	Number of PDUs to be sent																										
Entry Element	Size	Description																																		
Number	4	Number of PDUs to be sent																																		

Attribute Name	Id	M	Size (bytes)	Description		
				DataLength	1	Size of data packet to send
				DutyCycle	1	Average duty cycle (%). It must be the average because of the randomness of the CSMA/CA
				LCID	2	LCID of the broadcast data to be sent
				Priority	1	Priority of data packets to be sent
MACActionMgmtCon	0x006B	BN		Forces establishment/close of the management connection		
				Entry Element	Size	Description
				Node	6	EUI-48 of the Service Node
				Connect	1	0 – Close the management connection 1 – Open the management connection
MACActionMgmtMul	0x006C	BN		Forces establishment/close of the management multicast connection		
				Entry Element	Size	Description
				Node	6	EUI-48 of the Service Node

Attribute Name	Id	M	Size (bytes)	Description									
				<table border="1"> <tr> <td>Join</td> <td>1</td> <td> 0 – Leave the management multicast connection 1 – Join the management multicast connection </td> </tr> </table>	Join	1	0 – Leave the management multicast connection 1 – Join the management multicast connection						
Join	1	0 – Leave the management multicast connection 1 – Join the management multicast connection											
MACActionUnregisterBN	0x006D	BN	6	Trigger to unregister a given Service Node from the Subnetwork. PARAM: EUI-48 of the node being unregistered.									
MACActionConnCloseBN	0x006E	BN		Trigger to close an open connection.									
				<table border="1"> <thead> <tr> <th>Entry element</th> <th>Size</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>node</td> <td>6</td> <td>Eui48 of the node</td> </tr> <tr> <td>LCID</td> <td>2</td> <td>LCID of the connection to be closed.</td> </tr> </tbody> </table>	Entry element	Size	Description	node	6	Eui48 of the node	LCID	2	LCID of the connection to be closed.
Entry element	Size	Description											
node	6	Eui48 of the node											
LCID	2	LCID of the connection to be closed.											
MACActionSegmented4-32	0x006F	BN		<ul style="list-style-type: none"> • Trigger data transfer whit segmentation mechanism working (Convergence Layer) • Transmit PPDUs over established CL 4-32 Connection (with segmentation) • Trigger at least 1 packet segmented in at least 3 frames 									
				<table border="1"> <thead> <tr> <th>Entry Element</th> <th>Size</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Node</td> <td>6</td> <td>EUI-48 of the Service Node</td> </tr> <tr> <td>Length</td> <td>2</td> <td>Length of the data being transmitted (number or</td> </tr> </tbody> </table>	Entry Element	Size	Description	Node	6	EUI-48 of the Service Node	Length	2	Length of the data being transmitted (number or
Entry Element	Size	Description											
Node	6	EUI-48 of the Service Node											
Length	2	Length of the data being transmitted (number or											

Attribute Name	Id	M	Size (bytes)	Description		
						segments will depend on this length)
MACActionAppemuDataBurst	0x0080	BN		<p>Send a burst of data PDU-s with a test sequence using the Appemu connection to the node (if any)</p> <p>The data shall be transmitted with the flush bit to zero (0) when possible.</p>		
				Entry Element	Size	Description
				Node	6	EUI-48 of the Service Node
				Number	4	Number of PDU-s to be sent
				DataLength	1	Size of the data packets to be sent
				DutyCycle	1	Average duty cycle (percentage). It must be the average because of the randomness of the CSMA/CA
MACActionMgmtDataBurst	0x0081	M		<p>Send a burst of data PDU-s with a test sequence using the Management connection to the node (if any).</p> <p>The data shall be transmitted with the flush bit set to zero (0) when possible.</p>		
				Entry Element	Size	Description
				Node	6	EUI-48 of the Service Node, In service node this field will be ignored
				Number	4	Number of PDU-s to be sent

Attribute Name	Id	M	Size (bytes)	Description		
				DataLength	1	Size of the data packets to be sent
				DutyCycle	1	Average duty cycle (percentage). It must be the average because of the randomness of the CSMA/CA

4462 **6.2.4 Application PIB attributes**

4463 The following PIB attributes are used for general administration and maintenance of a OFDM PRIME
 4464 compliant device. These attributes do not affect the communication functionality, but enable easier
 4465 administration.

4466 These attributes shall be supported by both Base Node and Service Node devices.

4467 **Table 104 - 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.

Attribute Name	Size (in bits)	Id	Description						
AppListZCStatus		0x0078	<p>Zero Cross Status list. This list contains entry for each available zero cross detection circuits available in the system. Each element is sent together with reference time of zero cross close to frame beginning. If multiple entries are requested at the same time only first will be replied.</p> <table border="1" data-bbox="691 701 1461 1805"> <thead> <tr> <th data-bbox="691 701 906 797">Entry element</th> <th data-bbox="906 701 1058 797">Type</th> <th data-bbox="1058 701 1461 797">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="691 797 906 1805">ZCStatus</td> <td data-bbox="906 797 1058 1805">Byte</td> <td data-bbox="1058 797 1461 1805"> <p>Bit 7 : reserved, always 0</p> <p>Bits 5-6 : Terminal Block number</p> <p>0 : invalid</p> <p>1 : terminal block 1</p> <p>2 : terminal block 2</p> <p>3 : terminal block 3</p> <p>Bits 3-4 : Direction</p> <p>0 : unknown direction</p> <p>1 : falling</p> <p>2 : raising</p> <p>3 : reserved</p> <p>Bits 0-2: Status</p> <p>0 : available but unknown status</p> <p>1 : regular at 50Hz</p> <p>2 : regular at 60Hz</p> <p>3-5: reserved</p> <p>6: irregular intervals</p> <p>7: not available</p> </td> </tr> </tbody> </table>	Entry element	Type	Description	ZCStatus	Byte	<p>Bit 7 : reserved, always 0</p> <p>Bits 5-6 : Terminal Block number</p> <p>0 : invalid</p> <p>1 : terminal block 1</p> <p>2 : terminal block 2</p> <p>3 : terminal block 3</p> <p>Bits 3-4 : Direction</p> <p>0 : unknown direction</p> <p>1 : falling</p> <p>2 : raising</p> <p>3 : reserved</p> <p>Bits 0-2: Status</p> <p>0 : available but unknown status</p> <p>1 : regular at 50Hz</p> <p>2 : regular at 60Hz</p> <p>3-5: reserved</p> <p>6: irregular intervals</p> <p>7: not available</p>
Entry element	Type	Description							
ZCStatus	Byte	<p>Bit 7 : reserved, always 0</p> <p>Bits 5-6 : Terminal Block number</p> <p>0 : invalid</p> <p>1 : terminal block 1</p> <p>2 : terminal block 2</p> <p>3 : terminal block 3</p> <p>Bits 3-4 : Direction</p> <p>0 : unknown direction</p> <p>1 : falling</p> <p>2 : raising</p> <p>3 : reserved</p> <p>Bits 0-2: Status</p> <p>0 : available but unknown status</p> <p>1 : regular at 50Hz</p> <p>2 : regular at 60Hz</p> <p>3-5: reserved</p> <p>6: irregular intervals</p> <p>7: not available</p>							

4468 **6.3 Firmware upgrade**

4469 **6.3.1 General**

4470 The present section specifies firmware upgrade. Devices supporting PRIME may have several firmware inside
4471 them, at least one supporting the Application itself, and the one related to the PRIME protocol. Although it
4472 is possible that the application can perform the firmware upgrade of all the firmware images of the device,
4473 for instance DLMS/COSEM image transfer, using COSEM image transfer object, supporting PRIME firmware
4474 upgrade is mandatory in order to process to PRIME firmware upgrade independently of the application.

4475 **6.3.2 Requirements and features**

4476 This section specifies the firmware upgrade application, which is unique and mandatory for Base Nodes and
4477 Service Nodes.

4478 The most important features of the Firmware Upgrade mechanism are listed below. See following chapters
4479 for more information. The FU mechanism:

- 4480 • Shall be a part of management plane and therefore use the NULL SCS, as specified in section 0
- 4481 • Is able to work in unicast (default mode) and multicast (optional mode). The control messages
4482 are always sent using unicast connections, whereas data can be transmitted using both unicast
4483 and multicast. No broadcast should be used to transmit data.
- 4484 • May change the data packet sizes according to the channel conditions. The packet size will not
4485 be changed during the download process.
- 4486 • Is able to request basic information to the Service Nodes at anytime, such as device model,
4487 firmware version and FU protocol version.
- 4488 • Shall be abortable at anytime.
- 4489 • Shall check the integrity of the downloaded FW after completing the reception. In case of failure,
4490 the firmware upgrade application shall request a new retransmission.
- 4491 • The new firmware shall be executed in the Service Nodes only if they are commanded to do so.
4492 The FU application shall have to be able to set the moment when the reset takes place.
- 4493 • Must be able to reject the new firmware after a “test” period and switch to the old version. The
4494 duration of this test period has to be fixed by the FU mechanism.

4495 **6.3.3 General Description**

4496 **6.3.3.1 General**

4497 The Firmware Upgrade mechanism is able to work in unicast and multicast modes. All control messages are
4498 sent using unicast connections, whereas the data can be sent via unicast (by default) or multicast (only if
4499 supported by the manufacturer). Note that in order to ensure correct reception of the FW when Service
4500 Nodes from different vendors are upgraded, data packets shall not be sent via broadcast. Only unicast and
4501 multicast are allowed. A Node will reply only to messages sent via unicast. See chapter 6.3.5 for a detailed
4502 description of the control and information messages used by the FU mechanism.

4503 The unicast and multicast connections are set up by the Base Node. In case of supporting multicast, the Base
4504 Node shall request the Nodes from a specific vendor to join a specific multicast group, which is exclusively
4505 created to perform the firmware upgrade and is removed after finishing it.

4506 As said before, it is up to the vendor to use unicast or multicast for transmitting the data. In case of unicast
4507 data transmission, please note that the use of ARQ is an optional feature. Some examples showing the traffic
4508 between the Base Node and the Service Nodes in unicast and multicast are provided in 6.3.5.4.

4509 After completing the firmware download, each Service Node is committed by the Base Node to perform an
4510 integrity check on it. The firmware download will be restarted if the firmware image results to be corrupt. In
4511 other case, the Service Nodes will wait until they are commanded by the Base Node to execute the new
4512 firmware.

4513 The FU mechanism can setup the instant when the recently downloaded firmware is executed on the Service
4514 Nodes. Thus, the Base Node can choose to restart all Nodes at the same time or in several steps. After restart,
4515 each Service Node runs the new firmware for a time period specified by the FU mechanism. If this period
4516 expires without receiving any confirmation from the Base Node, or the Base Node decides to abort the
4517 upgrade process, the Service Nodes will reject the new firmware and switch to the old version. In any other
4518 case (a confirmation message is received) the Service Nodes will consider the new firmware as the only valid
4519 version and delete the old one.

4520 This is done in order to leave an “open back-door” in case that the new firmware is defect or corrupt. Please
4521 note that the Service Nodes are not allowed to discard any of the stored firmware versions until the final
4522 confirmation from the Base Node arrives or until the safety time period expires. The two last firmware
4523 upgrade steps explained above are shown in 6.3.5. See chapter 6.3.5.3 for a detailed description of the
4524 control messages.

Firmware Upgrade (version 1 -> version 2)

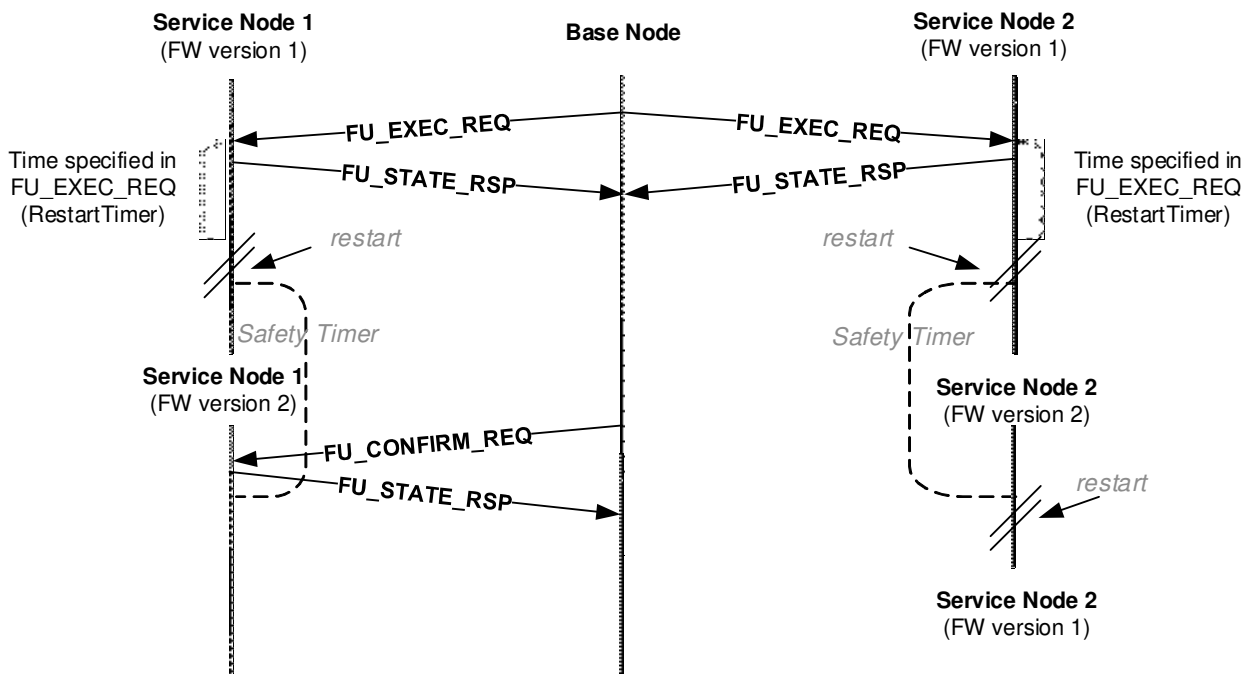


Figure 123 - Restarting de nodes and running the new firmware

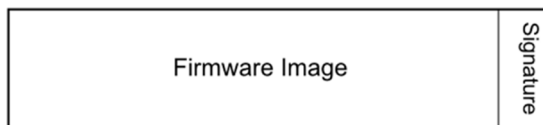
4525

4526

4527 **Note:** In normal circumstances, both Service Nodes should either accept or reject the new firmware version.
 4528 Both possibilities are shown above simultaneously for academic purposes.

4529 **6.3.3.2 Signed firmware**

4530 The "signed firmware" refers to the concatenation of the Firmware Image and the signature as shown in the
 4531 Figure 124. For now on in the document will be refered as signed firmware.



4532

4533

Figure 124 - Signed firmware diagram

4534 The payload transmitted in the Firmware Upgrade process shall be the signed firmware. For the SN to be able
 4535 to differentiate both, the signature will have a length defined in the FU_INIT_REQ's "Signature length" field.

4536 **6.3.3.3 Segmentation**

4537 The signed image is the information to be transferred, in order to process a firmware upgrade. The size of
 4538 the signed image will be called "ImageSize", and is measured in bytes. This image is divided in smaller
 4539 elements called pages that are easier to be transferred in packets. The "PageSize" may be one of the
 4540 following: 32 bytes, 64 bytes, 128 bytes or 192 bytes. This implies that the number of pages in a signed image
 4541 is calculated by the following formula:

4542

$$PageCount = \left\lceil \frac{ImageSize}{PageSize} \right\rceil + 1$$

4543

Every page will have a size specified by *PageSize*, except the last one that will contain the remaining bytes up to *ImageSize*.

4544

4545

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.

4546

4547

6.3.4 Firmware upgrade PIB attributes

4548

The following PIB attributes shall be supported by Service Nodes to support the firmware download application.

4549

4550

Table 105 - FU PIB attributes

Attribute Name	Size (in bits)	Id	Description
AppFwdlRunning	16	0x0070	Indicate 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.

4551

6.3.5 State machine

4552

6.3.5.1 General

4553

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.

4554

4555

4556

Table 106 - FU State Machine

	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>

	Description	Event	Output (or action to be performed)	Next state
		Receive FU_MISS_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>
		Any exception		<i>Exception</i>
Receiving	The FU application is receiving the Signed firmware.	Complete FW received, CRC OK and Signature OK		<i>Complete</i>
		Complete FW received and CRC not Ok or signature not OK		<i>Exception</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_INIT_REQ	FU_STATE_RSP (.State = 1)	<i>Receiving</i>
		Receive FU_DATA	(receiving data, normal behavior)	<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>
		Any exception		<i>Exception</i>

	Description	Event	Output (or action to be performed)	Next state
Complete	Upgrade completed, image integrity ok, the SN 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_INIT_REQ	FU_STATE_RSP (.State = 2)	<i>Complete</i>
		Receive FU_DATA	(ignore)	<i>Complete</i>
		Receive FU_EXEC_REQ with <i>RestartTimer</i> != 0	FU_STATE_RSP (.State = 3)	<i>Countdown</i>
		Receive FU_EXEC_REQ with <i>RestartTimer</i> = 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>
		Any exception		<i>Exception</i>
Countdown	Waiting until <i>RestartTimer</i> expires.	<i>RestartTimer</i> 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_INIT_REQ	FU_STATE_RSP (.State = 3)	<i>Countdown</i>
		Receive FU_DATA	(ignore)	<i>Countdown</i>
		Receive FU_EXEC_REQ with <i>RestartTimer</i> != 0	FU_STATE_RSP (.State = 3); (update <i>RestartTimer</i> and <i>SafetyTimer</i>)	<i>Countdown</i>

	Description	Event	Output (or action to be performed)	Next state
		Receive FU_EXEC_REQ with <i>RestartTimer</i> = 0	FU_STATE_RSP (.State = 4); (update <i>RestartTimer</i> and <i>SafetyTimer</i>)	<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>
		Any exception		<i>Exception</i>
Upgrade	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 = 4); (switch to <i>Exception</i> , FW rejected)	<i>Exception</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_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 = 4)	<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>
		Any exception		<i>Exception</i>
Exception	Upon any exception on the firmware upgrade service node will go into this state	Receive FU_INFO_REQ	FU_INFO_RSP	<i>Exception</i>
		Receive FU_STATE_REQ	FU_STATE_RSP (.State = 5)	<i>Exception</i>
		Receive FU_MISS_REQ	FU_STATE_RSP (.State = 5)	<i>Exception</i>
		Receive FU_INIT_REQ	FU_STATE_RSP (.State = 5)	<i>Exception</i>

4562 **6.3.5.2 State description**

4563 **6.3.5.2.1 Idle**

4564 The Service Nodes are in “Idle” state when they are not performing a firmware upgrade. The reception of a
4565 FU_INIT_REQ message is the only event that forces the Service Node to switch to the next state (“Receiving”).
4566 FU_KILL_REQ aborts the upgrade process and forces the Service Nodes to switch from any state to “Idle”.

4567 **6.3.5.2.2 Receiving**

4568 The Service Nodes receive the signed firmware via FU_DATA messages. Service Nodes report complete
4569 reception of the image answering with either an empty FU_MISS_LIST or an empty FU_MISS_BITMAP to the
4570 FU_MISS_REQ requests sent by the BN.

4571 If during the reception of the signed firmware the Service Node receives a block with a length that differs
4572 from the one configured in FU_INIT_REQ or a with an packet index out of bounds it should switch to
4573 “Exception” state with “Protocol” code.

4574 Once the download is complete, a Service Node shall check the integrity of the signed firmware by CRC
4575 calculation. If the CRC is wrong, the SN shall drop the signed firmware and switch to “Exception” state with
4576 “CRC verification fail” exception code.

4577 If the CRC results to be ok, the SN shall verify that the signed image is correctly signed with the manufacturer’s
4578 key. In case this verification fails, the SN shall drop the signed firmware and switch to “Exception” state with
4579 “Signature verification fail” exception code.

4580 If the signature is verified successfully, the SN shall switch to “Complete” state.

4581 The CRC check on the complete signed firmware and the later signature verification is mandatory, and is
4582 automatically started by the SNs. The service node shall not accept any image that is not properly signed.

4583 Note that these checks at SN side are not immediate. There may be a not negligible time interval between
4584 the message sent by the SN reporting that the reception is complete and the transition to “Complete”.

4585 **6.3.5.2.3 Complete**

4586 A Service Node in “Complete” state waits until reception of a FU_EXEC_REQ message. The Service Node may
4587 switch either to “Countdown” or “Upgrade” depending on the field *RestartTimer*, which specifies in which
4588 instant the Service Node has to reboot using the new firmware. If *RestartTimer* = 0, the Service Node
4589 immediately switches to “Upgrade”; else, the Service Node switches to “Countdown”.

4590 **6.3.5.2.4 Countdown**

4591 A Service Node in “Countdown” state waits a period of time specified in the *RestartTimer* field of a previous
4592 FU_EXEC_REQ message. When this timer expires, it automatically switches to “Upgrade”.

4593 FU_EXEC_REQ can be used in “Countdown” state to reset *RestartTimer* and *SafetyTimer*. In this case, both
 4594 timers have to be specified in FU_EXEC_REQ because both will be overwritten. Note that it is possible to force
 4595 the Node to immediately switch from “Countdown” to “Upgrade” state setting *RestartTimer* to zero.

4596 **6.3.5.2.5 Upgrade**

4597 A Service Node in “Upgrade” state shall run the new firmware during a time period specified in
 4598 FU_EXEC_REQ.SafetyTimer.

4599 If it does not receive any confirmation at all before this timer expires, the Service Node discards the new FW,
 4600 reboots with the old version and switches to “Exception” state with “Safety time expired” code.

4601 In case the SN receives a FU_KILL_REQ message it will discard the new FW, reboot with the old version and
 4602 switch to “Idle” state.

4603 **6.3.5.2.6 Exception**

4604 A Service Node can enter in exception state from any other state upon an event related to the Firmware
 4605 Upgrade that shall be notified to the Base Node as an exception.

4606 In case the SN receives a FU_KILL_REQ in “Exception” state it shall discard any ongoing FW upgrade progress
 4607 and switch to “idle” state. On any other event the SN will take no action and respond a FU_STATE_RSP to any
 4608 request with the code describing the specific exception. Exception state has a code, that shall have
 4609 information that can give more information on the exception happened. This code shall set the “temporary”
 4610 flag in case restarting the same Firmware Upgrade process could turn in success.

4611 There is a field up to the manufacturer of one byte for additional information about the exception, the format
 4612 of this field is out of the scope of this specification.

4613 **6.3.5.3 Control packets**

4614 **6.3.5.3.1 FU_INIT_REQ**

4615 The Base Node sends this packet in order to configure a Service Node for the Firmware Upgrade. If the Service
 4616 Node is in “Idle” state, it will change its state from “Idle” to “Receiving” and will answer with FU_STATE_RSP.
 4617 In any other case it will just answer sending FU_STATE_RSP.

4618 The content of FU_INIT_REQ is shown below.

4619 **Table 107 - Fields of FU_INIT_REQ**

Field	Length	Description
Type	4 bits	0 = FU_INIT_REQ.

Field	Length	Description
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 signed firmware in bytes.
CRC	32 bits	CRC of the signed firmware. 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.
Signature algorithm	4 bits	0 – no signature (not recommended to use in field, for testing purposes only) 1 – RSA 3072 + SHA-256 2 – ECDSA 256 + SHA-256 3-15 – Reserved for future use
Reserved	4 bits	Shall be 0 for this version of the document. Reserved for future use.
Signature length	8 bits	Length of the signature part of the signed firmware in bytes.

4620 **6.3.5.3.2 FU_EXEC_REQ**

4621 This packet is used by the Base Node to command a Service Node in “Complete” state to restart using the
4622 new firmware, once the complete image has been received by the Service Node. FU_EXEC_REQ specifies

4623 when the Service Node has to restart and how long the “safety” period shall be, as explained in 6.3.5.2.5.
 4624 Additionally, FU_EXEC_REQ can be used in “Countdown” state to reset the restart and the safety timers.

4625 Depending on the value of *RestartTimer*, a Service Node in “Complete” state may change either to
 4626 “Countdown” or to “Upgrade” state. In any case, the Service Node answers with FU_STATE_RSP.

4627 In “Countdown” state, the Base Node can reset *RestartTimer* and *SafetyTimer* with a FU_EXEC_REQ message
 4628 (both timers must be specified in the message because both will be overwritten).

4629 The content of this packet is described below.

4630 **Table 108 - 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.
<i>Reserved</i>	2 bits	0 .
<i>RestartTimer</i>	16 bits	0..65536 seconds; time before restarting with new FW.
<i>SafetyTimer</i>	16 bits	0..65536 seconds; time to test the new FW. It starts when the “Upgrade” state is entered.

4631 **6.3.5.3.3 FU_CONFIRM_REQ**

4632 This packet is sent by the Base Node to a Service Node in “Upgrade” state to confirm the current FW. If the
 4633 Service Node receives this message, it discards the old FW version and switches to “Idle” state. The Service
 4634 Node answers with FU_STATE_RSP when receiving this message.

4635 In any other state, the Service Node answers with FU_STATE_RSP without performing any additional actions.

4636 This packet contains the fields described below.

4637 **Table 109 - Fields of FU_CONFIRM_REQ**

Field	Length	Description
Type	4 bits	2 = FU_CONFIRM_REQ.

Field	Length	Description
Version	2 bits	0 for this version of the protocol.
<i>Reserved</i>	2 bits	0.

4638 6.3.5.3.4 FU_STATE_REQ

4639 This packet is sent by the Base Node in order to get the Firmware Upgrade state of a Service Node. The
4640 Service Node will answer with FU_STATE_RSP.

4641 This packet contains the fields described below.

4642 **Table 110 - 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	0.

4643 6.3.5.3.5 FU_KILL_REQ

4644 The Base Node sends this message to terminate the Firmware Upgrade process. A Service Node receiving this
4645 message will automatically switch to "Idle" state and optionally delete the downloaded data. The Service
4646 Node replies sending FU_STATE_RSP.

4647 The content of this packet is described below.

4648 **Table 111 - 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.

Field	Length	Description
<i>Reserved</i>	2 bits	0.

4649 **6.3.5.3.6 FU_STATE_RSP**

4650 **6.3.5.3.6.1 General**

4651 This packet is sent by the Service Node as an answer to FU_STATE_REQ, FU_KILL_REQ, FU_EXEC_REQ,
 4652 FU_CONFIRM_REQ or FU_INIT_REQ messages received through the unicast connection. It is used to notify
 4653 the Firmware Upgrade state in a Service Node.

4654 Additionally, FU_STATE_RSP is used as default response to all events that happen in states where they are
 4655 not foreseen (e.g. FU_EXEC_REQ in “Receiving” state, FU_INIT_REQ in “Upgrade” ...).

4656 This packet contains the fields described below.

4657 **Table 112 - 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	0.
State	4 bits	0 for Idle; 1 for Receiving; 2 for Complete; 3 for Countdown; 4 for Upgrade; 5 for Exception 6 to 15 reserved for future use.
Reserved	4 bits	0.

Field	Length	Description
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 State is Receiving).
Exception code	16 bits	Exception code describing the exception occurred (this field shall only be present if State is Exception) This field is described with detail in section 6.3.5.3.6.2

4658 **6.3.5.3.6.2 Exception code**

4659 The exception code has a number of fields that give more information to the Base Node about the exception
4660 that have happened during the Firmware Upgrade process.

4661 **Table 113 - Fields of Exception code**

Field	Length	Description
Permanent	1 bit	Flag used to inform if a retry on the firmware upgrade will not success, because the exception being permanent. 0 if the exception is temporary 1 if the exception is permanent
Code	7 bits	Code describing the type of exception that happened 0 – General: for an exception that do not fit any of the other codes 1 – Protocol: Page number out of bounds, page length mismatch from FU_INIT_REQ... 2 – CRC verification fail: If the CRC verification failed 3 – Invalid image: If the image was not a firmware image or not for the device 4 – Signature verification fail: the signature verification failed. 5 – Safety time expired: the safety time in "upgrade" state expires

Field	Length	Description
Manufacturer code	8 bits	Field that provides additional detail about the exception. This code is up to the manufacturer.

4662 **6.3.5.3.7 FU_DATA**

4663 This packet is sent by the Base Node to transfer a page of the signed firmware to a Service Node. No answer
4664 is expected by the Base Node.

4665 This packet contains the fields described below.

4666 **Table 114 - 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	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.

4667 **6.3.5.3.8 FU_MISS_REQ**

4668 This packet is sent by the Base Node to a Service Node to request information about the pages that are still
4669 to be received.

4670 If the Service Node is in “*Receiving*” state it will answer with a FU_MISS_BITMAP or FU_MISS_LIST message.

4671 If the Service Node is in any other state it will answer with a FU_STATE_RSP.

4672 This packet contains the fields described below.

4673

Table 115 - 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	0.
PageIndex	32 bits	Starting point to gather information about missing pages.

4674 **6.3.5.3.9 FU_MISS_BITMAP**

4675 This packet is sent by the Service Node as an answer to a FU_MISS_REQ. It carries the information about the
4676 pages that are still to be received.

4677 This packet will contain the fields described below.

4678

Table 116 - 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	0.
<i>Received</i>	32bits	Number of received pages.
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).

Field	Length	Description
Bitmap	<i>Variable</i>	<p>This bitmap contains the information about the status of each page.</p> <p>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+1</i> and so on.</p> <p>A '1' represents that a page is missing, a '0' represents that the page is already received.</p> <p>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.</p> <p>The maximum length of this field is <i>PageSize</i> bytes.</p>

4679 It is up to the Service Node to decide to send this type of packet or a FU_MISS_LIST message. It is usually
 4680 more efficient to transmit this kind of packets when the number of missing packets is not very low. But it is
 4681 up to the implementation to transmit one type of packet or the other. The Base Node should understand
 4682 both.

4683 In case a Service Node receives a FU_MISS_REQ during CRC calculation, it shall respond either with an empty
 4684 FU_MISS_BITMAP or an empty FU_MISS_LIST.

4685 **6.3.5.3.10 FU_MISS_LIST**

4686 This packet is sent by the Service Node as an answer to a FU_MISS_REQ. It carries the information about the
 4687 pages that are still to be received.

4688 This packet will contain the fields described below.

4689 **Table 117 - 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	0.
<i>Received</i>	32 bits	Number of received pages.

Field	Length	Description
PageIndexList	Variable	<p>List of pages that are still to be received. Each page is represented by its PageIndex, coded as a 32 bit integer.</p> <p>These pages should be sorted in ascending order (low to high), being possible to overflow to the PageIndex equal to zero to continue from the beginning.</p> <p>The first page index should be the same as the PageIndex 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).</p> <p>The maximum length of this field is PageSize bytes.</p>

4690 It is up to the Service Node to decide to transmit this packet type or a FU_MISS_BITMAP message. It is usually
 4691 more efficient to transmit this kind of packets when the missing packets are very sparse, but it is
 4692 implementation-dependent to transmit one type of packet or the other. The Base Node should understand
 4693 both.

4694 In case a Service Node receives a FU_MISS_REQ during CRC calculation, it shall respond either with an empty
 4695 FU_MISS_BITMAP or an empty FU_MISS_LIST.

4696 **6.3.5.3.11 FU_INFO_REQ**

4697 This packet is sent by a Base Node to request information from a Service Node, such as manufacturer, device
 4698 model, firmware version and other parameters specified by the manufacturer. The Service Node will answer
 4699 with one or more FU_INFO_RSP packets.

4700 This packet contains the fields described below.

4701 **Table 118 - 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.
Reserved	2 bits	0.

Field	Length	Description
InfoldList	Variable	List of identifiers with the information to retrieve. Each identifier is 1 byte long. The maximum length of this field is 32 bytes.

4702 The following identifiers are defined:

4703 **Table 119 - Infold possible values**

Infold	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.

4704 **6.3.5.3.12 FU_INFO_RSP**

4705 This packet is sent by a Service Node as a response to a FU_INFO_REQ message from the Base Node. A Service
4706 Node may have to send more than one FU_INFO_RSP when replying to a information request by the Base
4707 Node.

4708 This packet contains the fields described below.

4709 **Table 120 - 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	0.

Field	Length	Description
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).

4710 The InfoData field can contain several entries, the format of each entry is specified below.

4711 **Table 121 - Fields of each entry of InfoData in FU_INFO_RSP**

Field	Length	Description
Infold	8 bits	Identifier of the information as specified in 6.3.5.3.11.
<i>Reserved</i>	3 bits	0.
Length	5 bits	Length of the Data field (If Length is 0 it means that the specified Infold 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 Infold field. No value may be longer than 30 bytes.

4712 **6.3.5.4 Firmware integrity and authentication**

4713 **6.3.5.4.1 General**

4714 The firmware integrity and authentication is ensured by two means: the firmware CRC and the firmware
4715 signature. Both CRC and signature verifications are performed in the state “Receiving”, on the complete
4716 image after the receiving completion.

4717 **6.3.5.4.2 Image CRC**

4718 The role of the firmware upgrade CRC is to check the integrity of the image received from the Base Node,
4719 over the link Base Node – Service Node.

4720 The Base node, before initiating the firmware upgrade process, calculates a 32 bits CRC on the complete
4721 image, using the Generator polynomial $G(x)=x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^8+x^7+x^5+x^4+x^2+x+1$,
4722 and appends the result at the end of the firmware upgrade payload.

4723 After the receiving completion, the Service Node must verify the integrity of the whole image by the CRC
4724 recalculation. The firmware upgrade is deemed valid only when this CRC recalculation is successful.

4725 **6.3.5.4.3 Image signature**

4726 The role of the signature is to verify the integrity and the authenticity of the image as generated by the
4727 manufacturer. Indeed the firmware image, as generated by the image originator (the chip manufacturer) is
4728 provided to the Base Node via several different means. The firmware image signature provides evidence that
4729 the firmware image was not altered or substituted along all these transportation means and locations, and
4730 evidence that the concerned manufacturer is the originator.

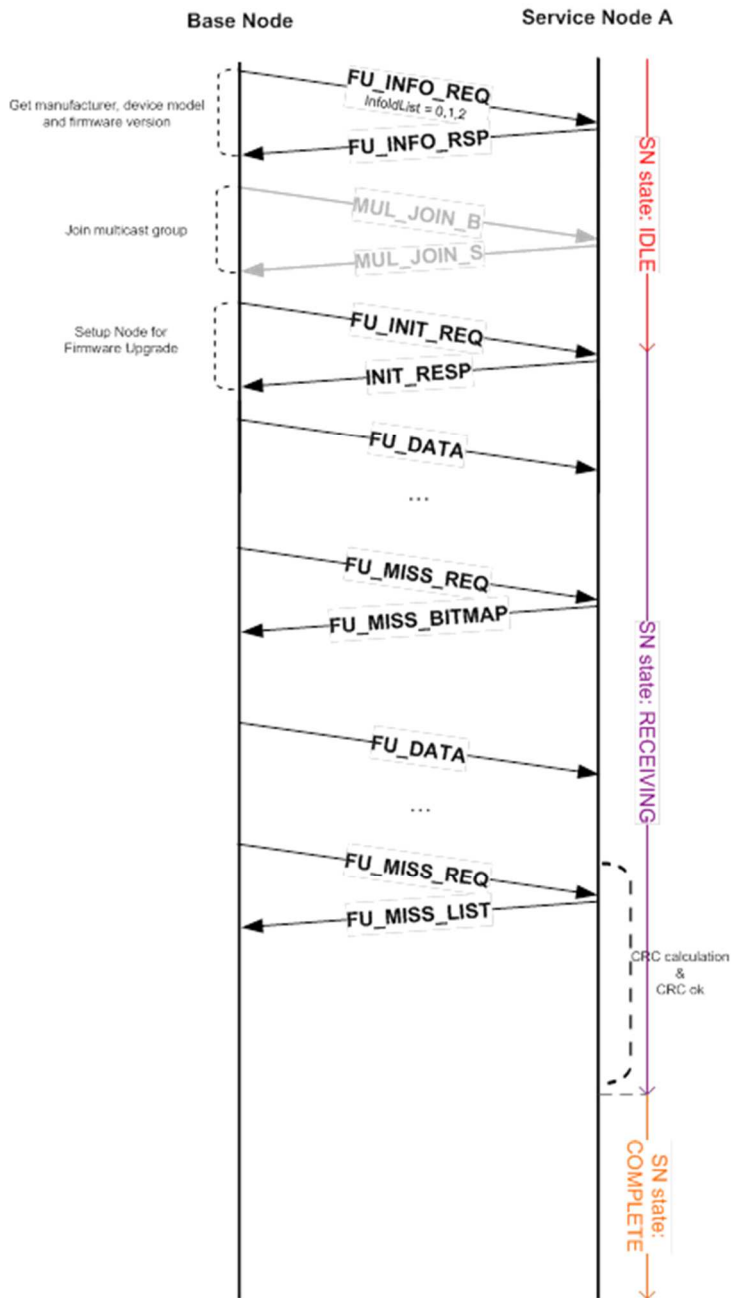
4731 The signature is generated by the originator of the firmware image, on the whole image using asymmetric
4732 key cryptography, and then included in the signed firmware to be delivered. The algorithm used for this
4733 purpose, how to equip the Service Nodes with the public key, and the infrastructure for certificate
4734 management are the scope of the firmware image generator. The algorithm used must comply with FIPS 186-
4735 4 standard, <http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-4.pdf>, preferably ECDSA 256 bits or RSA
4736 3072 bits as an alternate solution.

4737 Firmware image originators are strongly recommended for equipping the entities having in charge the
4738 spreading of the firmware image with tools allowing them to process to the verification before the
4739 deployment.

4740 After receiving the signed firmware the Service node must verify the firmware image signature. The image is
4741 deemed valid only when the verification is successful.

4742 **6.3.6 Examples**

4743 The figures below are an example of the traffic generated between the Base Node and the Service Node
4744 during the Firmware Upgrade process.

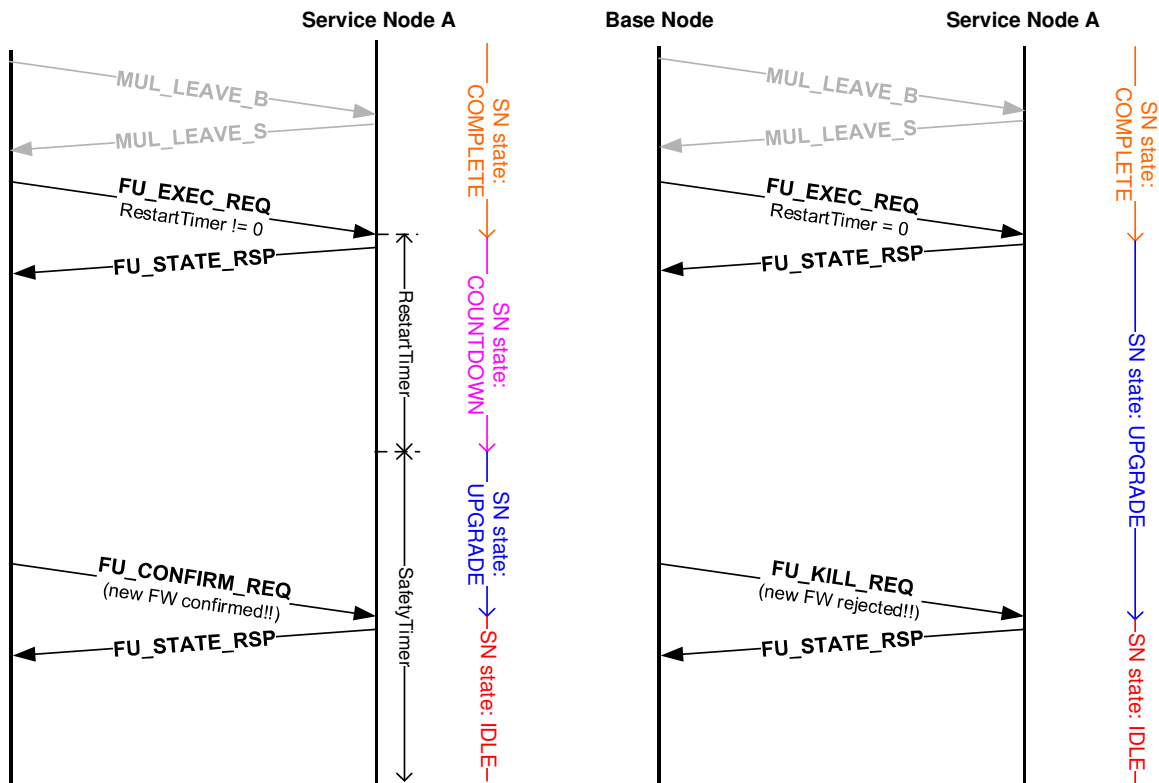


4745

4746

Figure 126 - Init Service Node and complete FW image

4747 Figure 126 shows the initialization of the process, the FW download and the integrity check of the image. In
 4748 the example above, the downloaded FW image is supposed to be complete before sending the last
 4749 FU_MISS_REQ. The Base Node sends it to verify its bitmap. In this example, FU_MISS_LIST has an empty
 4750 *PageIndexList* field, which means that the FW image is complete.



4751

4752

Figure 127 - Execute upgrade and confirm/reject new FW version

4753 Above it is shown how to proceed after completing the FW download. The Base Node commands the Service
 4754 Node to reboot either immediately (“Immediate Firmware Start”, *RestartTimer* = 0) or after a defined period
 4755 of time (“Delayed Firmware start”, *RestartTimer* != 0). After reboot, the Base Node can either confirm the
 4756 recently downloaded message sending a `FU_CONFIRM_REQ` or reject it (sending a `FU_KILL_REQ` or letting
 4757 the safety period expire doing nothing).

4758 6.4 Management interface description

4759 6.4.1 General


4760 Management functions defined in earlier sections shall be available over an abstract management interface
 4761 specified in this section. The management interface can be accessed over diverse media. Each physical media
 4762 shall specify its own management plane communication profile over which management information is
 4763 exchanged. It is mandatory for implementations to support PRIME management plane communication
 4764 profile. All other “management plane communication profiles” are optional and maybe mandated by certain
 4765 “application profiles” to use in specific cases.

4766 The present version of specifications describes two communication profiles, one of which is over this
 4767 specification NULL SCS and other over serial link.

4768 With these two communication profiles, it shall be possible to address the following use-cases:

- 4769 • Remote access of management interface over NULL SCS. This shall enable Base Node's use as a
- 4770 supervisory gateway for all devices in a Subnetwork
- 4771 • Local access of management interface (over peripherals like RS232, USBSerial etc) in a Service
- 4772 Node. Local access shall fulfill cases where a coprocessor exists for supervisory control of
- 4773 processor or when manual access is required over local physical interface for maintenance.

4774 Management data comprises of a 2 bytes header followed by payload information corresponding to the type
 4775 of information carried in message. The header comprises of a 10 bit length field and 6 bit message_id field.



4776
 4777 **Figure 128 - Management data frame**

4778 **Table 122 - 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 maybe redundant in some communication profiles (e.g. When transmitted over PRIME), but 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_id have standard semantics which should be respected by all PRIME 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 – Enhanced PIB Query 0x07 – Enhances PIB Response 0x08 to 0x0F: Reserved for future use. Vendors should not use these values for local purpose; 0x10 – 0x3F : Reserved for vendor specific use.


4779 **6.4.2 Payload format of management information**

4780 **6.4.2.1 Get PIB attribute query**

4781 This query is issued by a remote management entity that is interested in knowing values of PIB attributes
4782 maintained on a compliant device with this specification.

4783 The payload may comprise of a query on either a single PIB attribute or multiple attributes. For reasons of
4784 efficiency queries on multiple PIB attributes maybe aggregated in one single command. Given that the length
4785 of a PIB attribute identifier is constant, the number of attributes requested in a single command is derived
4786 from the overall MGMT.LEN field in header.

4787 The format of payload information is shown in the following figure.



4788
4789 **Figure 129 - Get PIB Attribute query. Payload**

4790 Fields of a GET request are summarized in table below:


4791 **Table 123 - 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 Attribute id. 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.

4792 **6.4.2.2 Get PIB attribute response**

4793 This data is sent out from a compliant device of this specification in response to a query of one or more PIB
4794 attributes. If a certain queried PIB attribute is not maintained on the device, it shall still respond to the query
4795 with value field containing all '1s' in the response.

4796 The format of payload is shown in the following figure.



4797

4798

Figure 130 - Get PIB Attribute response. Payload

4799

Fields of a GET request are summarized in table below:

4800

Table 124 - 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 Attribute id. 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 attribute. In case of a list attribute, value shall comprise of entire record corresponding to given index of PIB attribute
Next	1 byte	Index of next entry returned for corresponding PIB Attribute id. 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.

4801

Response to 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 if longer than the maximum segment size allowed to be carried over the NULL SCSS.

4802

4803

4804

6.4.2.3 Set PIB attribute

4805

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 SET command on several PIB attributes in one GPDU. The format of such an aggregated payload is shown in figure below:

4806

4807

4808




4809

4810

Figure 131 - Set PIB attribute. Aggregated payload

4811 For cases where the corresponding PIB attribute is only a trigger (all ACTION PIB attributes), there shall be
 4812 no associated value and the request data format shall be as shown below.



4813

4814

Figure 132 - Set PIB Attribute. ACTION PIB attributes

4815 It is assumed that the management entity sending out this information has already determined if the
 4816 corresponding attributes are supported at target device. This may be achieved by a previous query and its
 4817 response.

4818 **6.4.2.4 Reset statistics**

4819 This command has optional payload. In case there is no associated payload, the receiving device shall reset
 4820 all of its PIB statistical attributes.

4821 For cases when a remote management entity only intends to perform reset of selective PIB statistical
 4822 attributes, the payload shall contain a list of attributes that need to be reset. The format shall be the same
 4823 as shown in Section 6.4.2.1.

4824 Since there is no confirmation message going back from the device complying with this specification, the
 4825 management entity needs to send a follow-up PIB attribute query, in case it wants to confirm successful
 4826 completion of appropriate action.

4827 **6.4.2.5 Reboot device**

4828 There is no corresponding payload associated with this command. The command is complete in itself. The
 4829 receiving compliant device with this specification shall reboot itself on receipt of this message.

4830 It is mandatory for all implementations compliant with this specification to support this command and its
 4831 corresponding action.

4832 **6.4.2.6 Firmware upgrade**

4833 The payload in this case shall comprise of firmware upgrade commands and responses described in section
 4834 6.2.3.2 of the specification.

4835 **6.4.2.7 Enhanced PIB query**

4836 **6.4.2.7.1 General**

4837 This command let perform a variety of queries grouped in one. At the moment there is one query type that
 4838 can be performed, more queries will be added in future releases of the specification.

4839 The format of this command is shown in the Figure 133.

0x06	Query 1	Query 2	...	Query n
------	---------	---------	-----	---------

4840 **Figure 133 - Enhanced PIB query format**

1 bit	7 bits	8 bits
0	Iterator of 15 bits	

4841 **Figure 134 - Iterator short format**

1 bit	7 bits	<i>n</i> bytes
1	Iterator length (<i>n</i>)	Iterator of <i>n</i> bytes

4842 **Figure 135 - Iterator long format**

4843 The available queries are the ones listed in the Table 122.

4844 **6.4.2.7.2 PIB list query**

4845 This query is used to request the next elements in a collection of elements on a PIB element. Such as node
 4846 list or connection list.

4847 The format of this query is listed in Table 125

4848 **Table 125 - PIB List query format**

Element	Size(bytes)	Description
0x0E	1	Code for the PIB list query operation
Attribute ID	2	Attribute ID of the PIB
Number	1	Maximum number of records to retrieve

Element	Size(bytes)	Description
Iterator	*	Iterator returned in the last item if the PIB list response message, the response will start in the next element after that one. The constant value "0x0000" in this field will retrieve the first element

4849 **6.4.2.8 Enhanced PIB response**

4850 **6.4.2.8.1 General**

4851 This command let respond to a variety of queries requested in Enhanced PIB query. At the moment there is
4852 one response type, more response types will be added in future releases of the specification.

4853 The format of this command is shown in the Figure 136.

0x07	Response 1	Response 2	...	Response n
------	------------	------------	-----	------------

4854 **Figure 136 - Enhanced PIB response format**

4855 The available responses are listed in the Table 122.

4856 **6.4.2.8.2 PIB list response**

4857 This response is used to send information on lists of PIB collection elements, such as node list or connection
4858 list.

4859 The format of this command is shown in the Table 126.

4860 **Table 126 - PIB list response format**

Element	Size(bytes)	Description
0x0F	1	Code for the PIB list response operation
Attribute ID	2	Attribute ID of the PIB
Number	1	Number of records contained in this message
End of List	1	Length of every record

Element	Size(bytes)	Description
Iterator 1	*	Iterator of the record #1
Value 1	*	Value of the record #1
....
Iterator n	*	Iterator of the record #n
Value n	*	Value of the record #n

4861 The iterator has the same format as the ones described in section 6.4.2.7.2

4862 **6.4.3 NULL SSCS communication profile**

4863 This communication profile enables exchange of management information described in previous sections
 4864 over the NULL SSCS.

4865 The management entities at both transmitting and receiving ends are applications making use of the NULL
 4866 SSCS enumerated in Section 0 of this specs. Data is therefore exchanged as MAC Generic PDUs.

4867 **6.4.4 Serial communication profile**


4868 **6.4.4.1 Physical layer**

4869 The PHY layer maybe any serial link (e.g. RS232, USB Serial). The serial link is required to work 8N1
 4870 configuration at one of the following data rates:

4871 9600 bps, 19200 bps, 38400 bps, 57600 bps.

4872 **6.4.4.2 Data encapsulation for management messages**

4873 In order ensure robustness, the stream of data is encapsulated in HDLC-type frames which include a 2 byte
 4874 header and 4 byte CRC. All data is encapsulated between a starting flag-byte 0x7E and ending flag-byte 0x7E
 4875 as shown in Figure below:



4876

4877

Figure 137 - Data encapsulations for management messages

4878 If any of the intermediate data characters has the value 0x7E, it is preceded by an escape byte 0x7D, followed
 4879 by a byte derived from XORing the original character with byte 0x20. The same is done if there is a 0x7D
 4880 within the character stream. An example of such case is shown here:

4881 Msg to Tx: 0x01 0x02 0x7E 0x03 0x04 0x7D 0x05 0x06

4882 Actual Tx sequence: 0x01 0x02 0x7D 0x5E 0x03 0x04 0x7D 0x5D 0x05 0x06

4883 Escape Escape

4884 sequence sequence

4885 The 32 bit CRC at end of the frame covers both 'Header' and 'Payload' fields. The CRC is calculated over the
 4886 original data to be transmitted i.e. before byte stuffing of escape sequences described above is performed.
 4887 CRC calculation is

4888 The input polynomial $M(x)$ is formed as a polynomial whose coefficients are bits of the data being checked
 4889 (the first bit to check is the highest order coefficient and the last bit to check is the coefficient of order zero).
 4890 The Generator polynomial for the CRC is $G(x)=x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^8+x^7+x^5+x^4+x^2+x+1$.
 4891 The remainder $R(x)$ is calculated as the remainder from the division of $M(x)\cdot x^{32}$ by $G(x)$. The coefficients of
 4892 the remainder will then be the resulting CRC.

4893 6.4.5 TCP communication profile

4894 This communication profile enables exchange of management information described in previous sections
 4895 over a TCP socket. The socket number will be defined by the manufacturer and the device will have the role
 4896 of a TCP server.

4897 Management messages as described in Figure 128 are used in sequence with no extra header in this profile.
 4898 The message boundaries will be described with the length field of the management message.

4899 6.5 List of mandatory PIB attributes

4900 6.5.1 General

4901 PIB attributes listed in this section shall be supported by all implementations. PIB attributes that are not listed
 4902 in this section are optional and vendors may implement them at their choice. In addition to the PIB attributes,
 4903 the management command to reboot a certain device (as specified in 6.4.2.5) shall also be universally
 4904 supported.

4905 6.5.2 Mandatory PIB attributes common to all device types

4906 6.5.2.1 PHY PIB attribute

4907 (See Table 96)

4908 **Table 127 - PHY PIB common mandatory attributes**

Attribute Name	Id
phyStatsRxTotalCount	0x00A4
phyStatsBlkAvgEvm	0x00A5
phyEmaSmoothing	0x00A8

4909 **6.5.2.2 MAC PIB attributes**

4910 (See Table 97, Table 99 and Table 100)

4911 **Table 128 - MAC PIB comon mandatory attributes**

Attribute Name	Id
macEMASmoothing	0x0019
macCSMAR1 (non-Robust Modes only)	0x0034
macCSMAR2 (non-Robust Modes only)	0x0035
macCSMAR1Robust (Robust Modes supported)	0x003B
macCSMAR2Robust (Robust Modes supported)	0x003C

4912

Attribute Name	Id
MacCapabilities	0x002C

4913

List Attribute Name	Id
macListPhyComm	0x0059

4914 **6.5.2.3 Application PIB attributes**

4915 (See Table 104)

4916 **Table 129 - Applications PIB common mandatory attributes**

Attribute Name	Id
AppFwVersion	0x0075
AppVendorId	0x0076
AppProductId	0x0077

4917 **6.5.3 Mandatory Base Node attributes**4918 **6.5.3.1 MAC PIB attributes**

4919 (See Table 97 and Table 101)

4920 **Table 130 - MAC PIB Base Node mandatory attributes**

Attribute Name	Id
macBeaconsPerFrame	0x0013

4921

List Attribute Name	Id
macListRegDevices	0x0050
macListActiveConn	0x0051

4922 **6.5.4 Mandatory Service Node attributes**4923 **6.5.4.1 MAC PIB attributes**

4924 (See Table 99, Table 101 and Table 103).

4925 **Table 131 - 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

4926

List Attribute Name	Id
macListSwitchTable	0x0053
macListAvailableSwitches	0x0056

4927

Attribute Name	Id
MACActionTxData	0x0060
MACActionConnClose	0x0061
MACActionRegReject	0x0062
MACActionProReject	0x0063
MACActionUnregister	0x0064
macSecDUK	0x005B

4928 **6.5.4.2 Application PIB attributes**

4929 (See Table 105)

4930 **Table 132 - APP PIB Service Node mandatory attributes**

Attribute Name	Id
<u>AppFwdlRunning</u>	0x0070
<u>AppFwdlRxPktCount</u>	0x0071

4931
4932
4933

**Annex A
(informative)
Examples of CRC**

4934 **CRC-8 Example**

4935 The table below gives the CRC-8 examples (see section 3.4.3) calculated for several specified strings

4936 **Table 133 - Examples of CRC-8 calculated for various ASCII strings**

String	CRC-8
'T'	0xab
"THE"	0xa0
0x03, 0x73	0x61
0x01, 0x3f	0xa8
"123456789"	0xf4

4937 **CRC-32 Example**

4938 The table below gives the CRC-32 example (see section 3.4.3)

4939 **Table 134 - Example of CRC-32**

String	CRC-32
0x30 0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39 0x30 0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39 0x30 0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39 0x30 0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39 0x30 0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39	0x24a56cf5

4940
4941
4942

Annex B
(normative)
EVM calculation

4943 This annex describes calculation of the EVM by a reference receiver, assuming accurate synchronization and
4944 FFT window placement. Let

- 4945
- r_k^i denotes the FFT output for symbol i and k are the indices of data subcarriers.
 - $\Delta b_k \in \{0,1,\dots,P-1\}$ represents the decision on the received information symbol coded in the phase increment.
 - $P = 2, 4, \text{ or } 8$ in the case of DBPSK, DQPSK or D8PSK, respectively.
- 4949

4950 The EVM definition is then given by;

$$EVM = \frac{\sum_{i=1}^L \sum_{k \in \{\text{data subcarriers}\}} \left(\text{abs} \left(r_k^i - r_{k-1}^i e^{-j*2*\pi/P \times \Delta b_{k-1}} \right) \right)^2}{\sum_{i=1}^L \sum_{k \in \{\text{data subcarriers}\}} \left(\text{abs} \left(r_k^i \right) \right)^2}$$

4951

4952 In the above, $\text{abs}(\cdot)$ refers to the magnitude of a complex number. L is the number of OFDM symbols in the
4953 most recently received PPDU, over which the EVM is calculated.

4954 The noise can be estimated as the numerator of the EVM. The RSSI can be estimated as the denominator of
4955 the EVM. The SNR can be estimated as the reciprocal of the EVM above plus 3dB due to differential decoding.

4956
4957
4958
4959

**Annex C
(informative)
Interleaving matrixes (N_{CH} = 1)**

Table 135 - 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

4960

Table 136 - 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

4961

Table 137 - 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

4962

Table 138 - 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

4963

4964
4965
4966

Annex D (normative) MAC layer constants

4967 This section defines all the MAC layer constants.

4968

Table 139 - Table of MAC constants

Constant	Value	Description
MACBeaconLength1	5 symbols	Length of beacon in symbols. Type A frame and DBPSK_CC modulation
MACBeaconLength2	16 symbols	Length of beacon in symbols. Type B frame and DQPSK_RC modulation
MACBeaconLength3	24 symbols	Length of beacon in symbols. Type B frame and DBPSK_RC modulation
MACBeaconLength4	19 symbols	Length of beacon in symbols. Type BC frame and DQPSK_R modulation
MACBeaconLength5	27 symbols	Length of beacon in symbols. Type BC frame and DBPSK_R modulation
MACMinSCPLength	64 symbols	Minimum length of SCP.
MACPriorityLevels	4	Number of levels of priority supported by the system.
MACMinRobustnessLevel	DBPSK_CC	Weakest modulation scheme
MACCtrlPktPriority	1	MAC Priority used to transmit all the Control Packets except ALV Control Packets
MACALVCtrlTketPriority	0	MAC Priority used to transmit ALV control Packets
MACSuperFrameLength	32	Number of frames that defines the superframe

Constant	Value	Description
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.
MACConcurrentAliveProcedure	2	The number of Alive procedure a Service Node shall support at any given time
$N_{\text{miss-beacon}}$	5	Number of superframes a Service Node does not receive an expected beacon before considering its Switch Node as unavailable.
ARQMaxTxCount	32	Maximum allowed retransmission count before an ARQ connection must be closed
ARQCongClrTime	10 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.

4969
4970
4971

**Annex E
(normative)
Convergence layer constants**

4972 The following TYPE values are defined for use by Convergence layers from chapter 5.

4973

Table 140 - 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_DATA	6

4974 The following LCID values apply for broadcast connections defined by Convergence layers from chapter 5.

4975

Table 141 - LCID value assignments

LCID Symbolic Name	Value	MAC Scope
LCI_CL_IPv4_BROADCAST	1	Broadcast.
LCI_CL_432_BROADCAST	2	Broadcast.

4976 The following Result values are defined for Convergence layer primitives.

4977

Table 142 - 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.

Result	Description
Timeout = 2	A timed out occurs during the SCS service processing
Not Registered = 6	The service node is not currently registered to a Subnetwork.
Unsupported SP = 14	Device doesn't support SP > 0 but asked for an encrypted connection establishment

4978
4979
4980

Annex F (normative) Profiles

4981 Given the different applications which are foreseen for this specification compliant products, it is necessary
4982 to define different profiles. Profiles cover the functionalities that represent the respective feature set. They
4983 need to be implemented as written in order to assure interoperability.

4984 This specification has a number of options, which, if exercised in different ways by different vendors, will
4985 hamper both compliance testing activities and future product interoperability. The profiles further restrict
4986 those options so as to promote interoperability and testability.

4987 A specific profile will dictate which capabilities a Node negotiates through the Registering and Promotion
4988 processes.

4989 **F.1 Smart Metering Profile**

4990 The following options will be either mandatory or optional for Smart Metering Nodes.

4991 REG.CAP_SW:

- 4992 • Base Node: Set to 1.
- 4993 • Service Node: Set to 1.

4994 REG.CAP_PA:

- 4995 • Base Node: optional.
- 4996 • Service Node: optional.

4997 REG.CAP_CFP:

- 4998 • Base Node: optional.
- 4999 • Service Node: optional.

5000 REG.CAP_DC

- 5001 • Base Node: optional.
- 5002 • Service Node: optional.

5003 REG.CAP_MC

- 5004 • Base Node: Set to 1.
- 5005 • Service Node: optional.

5006 REG.CAP_RM

- 5007 • Base Node: Set to 1.
- 5008 • Service Node: Set to 1.

- 5009 REG.CAP_ARQ
- 5010 • Base Node: optional.
- 5011 • Service Node: optional.
- 5012 PRO.SWC_DC
- 5013 • Service Node: optional.
- 5014 PRO.SWC_MC
- 5015 • Service Node: optional.
- 5016 PRO.SWC_RM
- 5017 • Service Node: Set to 1.
- 5018 PRO.SWC_ARQ
- 5019 • Service Node: optional.

5020
5021
5022

**Annex G
(informative)
List of frequencies used**

5023 The tables below give the exact center frequencies (in Hz) for the 97 subcarriers of the OFDM signal, channel
5024 by channel.

5025 Note that a guard period of 15 subcarriers is kept between any two consecutive channels.

5026

Table 143 - Channel 1: List of frequencies used

#	Frequency	#	Frequency	#	Frequency	#	Frequency
86	41992.18750	111	54199.21875	136	66406.25000	161	78613.28125
87	42480.46875	112	54687.50000	137	66894.53125	162	79101.56250
88	42968.75000	113	55175.78125	138	67382.81250	163	79589.84375
89	43457.03125	114	55664.06250	139	67871.09375	164	80078.12500
90	43945.31250	115	56152.34375	140	68359.37500	165	80566.40625
91	44433.59375	116	56640.62500	141	68847.65625	166	81054.68750
92	44921.87500	117	57128.90625	142	69335.93750	167	81542.96875
93	45410.15625	118	57617.18750	143	69824.21875	168	82031.25000
94	45898.43750	119	58105.46875	144	70312.50000	169	82519.53125
95	46386.71875	120	58593.75000	145	70800.78125	170	83007.81250
96	46875.00000	121	59082.03125	146	71289.06250	171	83496.09375
97	47363.28125	122	59570.31250	147	71777.34375	172	83984.37500
98	47851.56250	123	60058.59375	148	72265.62500	173	84472.65625

#	Frequency	#	Frequency	#	Frequency	#	Frequency
99	48339.84375	124	60546.87500	149	72753.90625	174	84960.93750
100	48828.12500	125	61035.15625	150	73242.18750	175	85449.21875
101	49316.40625	126	61523.43750	151	73730.46875	176	85937.50000
102	49804.68750	127	62011.71875	152	74218.75000	177	86425.78125
103	50292.96875	128	62500.00000	153	74707.03125	178	86914.06250
104	50781.25000	129	62988.28125	154	75195.31250	179	87402.34375
105	51269.53125	130	63476.56250	155	75683.59375	180	87890.62500
106	51757.81250	131	63964.84375	156	76171.87500	181	88378.90625
107	52246.09375	132	64453.12500	157	76660.15625	182	88867.18750
108	52734.37500	133	64941.40625	158	77148.43750		
109	53222.65625	134	65429.68750	159	77636.71875		
110	53710.93750	135	65917.96875	160	78125.00000		

5027

Table 144 - Channel 2: List of frequencies used

#	Frequency	#	Frequency	#	Frequency	#	Frequency
198	96679.68750	223	108886.71875	248	121093.75000	273	133300.78125
199	97167.96875	224	109375.00000	249	121582.03125	274	133789.06250
200	97656.25000	225	109863.28125	250	122070.31250	275	134277.34375

#	Frequency	#	Frequency	#	Frequency	#	Frequency
201	98144.53125	226	110351.56250	251	122558.59375	276	134765.62500
202	98632.81250	227	110839.84375	252	123046.87500	277	135253.90625
203	99121.09375	228	111328.12500	253	123535.15625	278	135742.18750
204	99609.37500	229	111816.40625	254	124023.43750	279	136230.46875
205	100097.65625	230	112304.68750	255	124511.71875	280	136718.75000
206	100585.93750	231	112792.96875	256	125000.00000	281	137207.03125
207	101074.21875	232	113281.25000	257	125488.28125	282	137695.31250
208	101562.50000	233	113769.53125	258	125976.56250	283	138183.59375
209	102050.78125	234	114257.81250	259	126464.84375	284	138671.87500
210	102539.06250	235	114746.09375	260	126953.12500	285	139160.15625
211	103027.34375	236	115234.37500	261	127441.40625	286	139648.43750
212	103515.62500	237	115722.65625	262	127929.68750	287	140136.71875
213	104003.90625	238	116210.93750	263	128417.96875	288	140625.00000
214	104492.18750	239	116699.21875	264	128906.25000	289	141113.28125
215	104980.46875	240	117187.50000	265	129394.53125	290	141601.56250
216	105468.75000	241	117675.78125	266	129882.81250	291	142089.84375
217	105957.03125	242	118164.06250	267	130371.09375	292	142578.12500

#	Frequency	#	Frequency	#	Frequency	#	Frequency
218	106445.31250	243	118652.34375	268	130859.37500	293	143066.40625
219	106933.59375	244	119140.62500	269	131347.65625	294	143554.68750
220	107421.87500	245	119628.90625	270	131835.93750		
221	107910.15625	246	120117.18750	271	132324.21875		
222	108398.43750	247	120605.46875	272	132812.50000		

5028

Table 145 - Channel 3: List of frequencies used

#	Frequency	#	Frequency	#	Frequency	#	Frequency
310	151367.18750	335	163574.21875	360	175781.25000	385	187988.28125
311	151855.46875	336	164062.50000	361	176269.53125	386	188476.56250
312	152343.75000	337	164550.78125	362	176757.81250	387	188964.84375
313	152832.03125	338	165039.06250	363	177246.09375	388	189453.12500
314	153320.31250	339	165527.34375	364	177734.37500	389	189941.40625
315	153808.59375	340	166015.62500	365	178222.65625	390	190429.68750
316	154296.87500	341	166503.90625	366	178710.93750	391	190917.96875
317	154785.15625	342	166992.18750	367	179199.21875	392	191406.25000
318	155273.43750	343	167480.46875	368	179687.50000	393	191894.53125
319	155761.71875	344	167968.75000	369	180175.78125	394	192382.81250

#	Frequency	#	Frequency	#	Frequency	#	Frequency
320	156250.00000	345	168457.03125	370	180664.06250	395	192871.09375
321	156738.28125	346	168945.31250	371	181152.34375	396	193359.37500
322	157226.56250	347	169433.59375	372	181640.62500	397	193847.65625
323	157714.84375	348	169921.87500	373	182128.90625	398	194335.93750
324	158203.12500	349	170410.15625	374	182617.18750	399	194824.21875
325	158691.40625	350	170898.43750	375	183105.46875	400	195312.50000
326	159179.68750	351	171386.71875	376	183593.75000	401	195800.78125
327	159667.96875	352	171875.00000	377	184082.03125	402	196289.06250
328	160156.25000	353	172363.28125	378	184570.31250	403	196777.34375
329	160644.53125	354	172851.56250	379	185058.59375	404	197265.62500
330	161132.81250	355	173339.84375	380	185546.87500	405	197753.90625
331	161621.09375	356	173828.12500	381	186035.15625	406	198242.18750
332	162109.37500	357	174316.40625	382	186523.43750		
333	162597.65625	358	174804.68750	383	187011.71875		
334	163085.93750	359	175292.96875	384	187500.00000		

Table 146 - Channel 4: List of frequencies used

#	Frequency	#	Frequency	#	Frequency	#	Frequency
422	206054.68750	447	218261.71875	472	230468.75000	497	242675.78125
423	206542.96875	448	218750.00000	473	230957.03125	498	243164.06250
424	207031.25000	449	219238.28125	474	231445.31250	499	243652.34375
425	207519.53125	450	219726.56250	475	231933.59375	500	244140.62500
426	208007.81250	451	220214.84375	476	232421.87500	501	244628.90625
427	208496.09375	452	220703.12500	477	232910.15625	502	245117.18750
428	208984.37500	453	221191.40625	478	233398.43750	503	245605.46875
429	209472.65625	454	221679.68750	479	233886.71875	504	246093.75000
430	209960.93750	455	222167.96875	480	234375.00000	505	246582.03125
431	210449.21875	456	222656.25000	481	234863.28125	506	247070.31250
432	210937.50000	457	223144.53125	482	235351.56250	507	247558.59375
433	211425.78125	458	223632.81250	483	235839.84375	508	248046.87500
434	211914.06250	459	224121.09375	484	236328.12500	509	248535.15625
435	212402.34375	460	224609.37500	485	236816.40625	510	249023.43750
436	212890.62500	461	225097.65625	486	237304.68750	511	249511.71875
437	213378.90625	462	225585.93750	487	237792.96875	512	250000.00000
438	213867.18750	463	226074.21875	488	238281.25000	513	250488.28125

#	Frequency	#	Frequency	#	Frequency	#	Frequency
439	214355.46875	464	226562.50000	489	238769.53125	514	250976.56250
440	214843.75000	465	227050.78125	490	239257.81250	515	251464.84375
441	215332.03125	466	227539.06250	491	239746.09375	516	251953.12500
442	215820.31250	467	228027.34375	492	240234.37500	517	252441.40625
443	216308.59375	468	228515.62500	493	240722.65625	518	252929.68750
444	216796.87500	469	229003.90625	494	241210.93750		
445	217285.15625	470	229492.18750	495	241699.21875		
446	217773.43750	471	229980.46875	496	242187.50000		

5030

Table 147 - Channel 5: List of frequencies used

#	Frequency	#	Frequency	#	Frequency	#	Frequency
534	260742.18750	559	272949.21875	584	285156.25000	609	297363.28125
535	261230.46875	560	273437.50000	585	285644.53125	610	297851.56250
536	261718.75000	561	273925.78125	586	286132.81250	611	298339.84375
537	262207.03125	562	274414.06250	587	286621.09375	612	298828.12500
538	262695.31250	563	274902.34375	588	287109.37500	613	299316.40625
539	263183.59375	564	275390.62500	589	287597.65625	614	299804.68750
540	263671.87500	565	275878.90625	590	288085.93750	615	300292.96875

#	Frequency	#	Frequency	#	Frequency	#	Frequency
541	264160.15625	566	276367.18750	591	288574.21875	616	300781.25000
542	264648.43750	567	276855.46875	592	289062.50000	617	301269.53125
543	265136.71875	568	277343.75000	593	289550.78125	618	301757.81250
544	265625.00000	569	277832.03125	594	290039.06250	619	302246.09375
545	266113.28125	570	278320.31250	595	290527.34375	620	302734.37500
546	266601.56250	571	278808.59375	596	291015.62500	621	303222.65625
547	267089.84375	572	279296.87500	597	291503.90625	622	303710.93750
548	267578.12500	573	279785.15625	598	291992.18750	623	304199.21875
549	268066.40625	574	280273.43750	599	292480.46875	624	304687.50000
550	268554.68750	575	280761.71875	600	292968.75000	625	305175.78125
551	269042.96875	576	281250.00000	601	293457.03125	626	305664.06250
552	269531.25000	577	281738.28125	602	293945.31250	627	306152.34375
553	270019.53125	578	282226.56250	603	294433.59375	628	306640.62500
554	270507.81250	579	282714.84375	604	294921.87500	629	307128.90625
555	270996.09375	580	283203.12500	605	295410.15625	630	307617.18750
556	271484.37500	581	283691.40625	606	295898.43750		
557	271972.65625	582	284179.68750	607	296386.71875		

#	Frequency	#	Frequency	#	Frequency	#	Frequency
558	272460.93750	583	284667.96875	608	296875.00000		

5031

Table 148 - Channel 6: List of frequencies used

#	Frequency	#	Frequency	#	Frequency	#	Frequency
646	315429.68750	671	327636.71875	696	339843.75000	721	352050.78125
647	315917.96875	672	328125.00000	697	340332.03125	722	352539.06250
648	316406.25000	673	328613.28125	698	340820.31250	723	353027.34375
649	316894.53125	674	329101.56250	699	341308.59375	724	353515.62500
650	317382.81250	675	329589.84375	700	341796.87500	725	354003.90625
651	317871.09375	676	330078.12500	701	342285.15625	726	354492.18750
652	318359.37500	677	330566.40625	702	342773.43750	727	354980.46875
653	318847.65625	678	331054.68750	703	343261.71875	728	355468.75000
654	319335.93750	679	331542.96875	704	343750.00000	729	355957.03125
655	319824.21875	680	332031.25000	705	344238.28125	730	356445.31250
656	320312.50000	681	332519.53125	706	344726.56250	731	356933.59375
657	320800.78125	682	333007.81250	707	345214.84375	732	357421.87500
658	321289.06250	683	333496.09375	708	345703.12500	733	357910.15625
659	321777.34375	684	333984.37500	709	346191.40625	734	358398.43750

#	Frequency	#	Frequency	#	Frequency	#	Frequency
660	322265.62500	685	334472.65625	710	346679.68750	735	358886.71875
661	322753.90625	686	334960.93750	711	347167.96875	736	359375.00000
662	323242.18750	687	335449.21875	712	347656.25000	737	359863.28125
663	323730.46875	688	335937.50000	713	348144.53125	738	360351.56250
664	324218.75000	689	336425.78125	714	348632.81250	739	360839.84375
665	324707.03125	690	336914.06250	715	349121.09375	740	361328.12500
666	325195.31250	691	337402.34375	716	349609.37500	741	361816.40625
667	325683.59375	692	337890.62500	717	350097.65625	742	362304.68750
668	326171.87500	693	338378.90625	718	350585.93750		
669	326660.15625	694	338867.18750	719	351074.21875		
670	327148.43750	695	339355.46875	720	351562.50000		

5032

Table 149 - Channel 7: List of frequencies used

#	Frequency	#	Frequency	#	Frequency	#	Frequency
758	370117.18750	783	382324.21875	808	394531.25000	833	406738.28125
759	370605.46875	784	382812.50000	809	395019.53125	834	407226.56250
760	371093.75000	785	383300.78125	810	395507.81250	835	407714.84375
761	371582.03125	786	383789.06250	811	395996.09375	836	408203.12500

#	Frequency	#	Frequency	#	Frequency	#	Frequency
762	372070.31250	787	384277.34375	812	396484.37500	837	408691.40625
763	372558.59375	788	384765.62500	813	396972.65625	838	409179.68750
764	373046.87500	789	385253.90625	814	397460.93750	839	409667.96875
765	373535.15625	790	385742.18750	815	397949.21875	840	410156.25000
766	374023.43750	791	386230.46875	816	398437.50000	841	410644.53125
767	374511.71875	792	386718.75000	817	398925.78125	842	411132.81250
768	375000.00000	793	387207.03125	818	399414.06250	843	411621.09375
769	375488.28125	794	387695.31250	819	399902.34375	844	412109.37500
770	375976.56250	795	388183.59375	820	400390.62500	845	412597.65625
771	376464.84375	796	388671.87500	821	400878.90625	846	413085.93750
772	376953.12500	797	389160.15625	822	401367.18750	847	413574.21875
773	377441.40625	798	389648.43750	823	401855.46875	848	414062.50000
774	377929.68750	799	390136.71875	824	402343.75000	849	414550.78125
775	378417.96875	800	390625.00000	825	402832.03125	850	415039.06250
776	378906.25000	801	391113.28125	826	403320.31250	851	415527.34375
777	379394.53125	802	391601.56250	827	403808.59375	852	416015.62500
778	379882.81250	803	392089.84375	828	404296.87500	853	416503.90625

#	Frequency	#	Frequency	#	Frequency	#	Frequency
779	380371.09375	804	392578.12500	829	404785.15625	854	416992.18750
780	380859.37500	805	393066.40625	830	405273.43750		
781	381347.65625	806	393554.68750	831	405761.71875		
782	381835.93750	807	394042.96875	832	406250.00000		

5033

Table 150 - Channel 8: List of frequencies used

#	Frequency	#	Frequency	#	Frequency	#	Frequency
870	424804.68750	895	437011.71875	920	449218.75000	945	461425.78125
871	425292.96875	896	437500.00000	921	449707.03125	946	461914.06250
872	425781.25000	897	437988.28125	922	450195.31250	947	462402.34375
873	426269.53125	898	438476.56250	923	450683.59375	948	462890.62500
874	426757.81250	899	438964.84375	924	451171.87500	949	463378.90625
875	427246.09375	900	439453.12500	925	451660.15625	950	463867.18750
876	427734.37500	901	439941.40625	926	452148.43750	951	464355.46875
877	428222.65625	902	440429.68750	927	452636.71875	952	464843.75000
878	428710.93750	903	440917.96875	928	453125.00000	953	465332.03125
879	429199.21875	904	441406.25000	929	453613.28125	954	465820.31250
880	429687.50000	905	441894.53125	930	454101.56250	955	466308.59375

#	Frequency	#	Frequency	#	Frequency	#	Frequency
881	430175.78125	906	442382.81250	931	454589.84375	956	466796.87500
882	430664.06250	907	442871.09375	932	455078.12500	957	467285.15625
883	431152.34375	908	443359.37500	933	455566.40625	958	467773.43750
884	431640.62500	909	443847.65625	934	456054.68750	959	468261.71875
885	432128.90625	910	444335.93750	935	456542.96875	960	468750.00000
886	432617.18750	911	444824.21875	936	457031.25000	961	469238.28125
887	433105.46875	912	445312.50000	937	457519.53125	962	469726.56250
888	433593.75000	913	445800.78125	938	458007.81250	963	470214.84375
889	434082.03125	914	446289.06250	939	458496.09375	964	470703.12500
890	434570.31250	915	446777.34375	940	458984.37500	965	471191.40625
891	435058.59375	916	447265.62500	941	459472.65625	966	471679.68750
892	435546.87500	917	447753.90625	942	459960.93750		
893	436035.15625	918	448242.18750	943	460449.21875		
894	436523.43750	919	448730.46875	944	460937.50000		

5034
5035
5036

Annex H
(informative)
Informative

5037 **H.1 Data exchange between to IP communication peers**

5038 This example shows the primitive exchange between a service node (192.168.0.100/24) and a base node
5039 when the former wants to exchange IP packets with a third service node (192.168.0.101/24) whose IP address
5040 is in the same IP Subnetwork.

5041 This example makes the following assumptions:

- 5042 • Service node (192.168.0.100) IPv4 SSCS does not exist so it needs to start a IPv4 SSCS and register
5043 its IP address in the base node prior to the exchange of IP packets.
5044 • Service node (192.168.0.101) has already registered its IP Address in the base node.

5045 The steps illustrated in next page are:

5046 1. The IPv4 layer of the service node (192.168.0.100) invokes the CL_IPv4_ESTABLISH.request primitive. To
5047 establish IPv4 SSCS, it is required,

5048 a. To establish a connection with the base node so all address resolution messages can be exchanged
5049 over it.

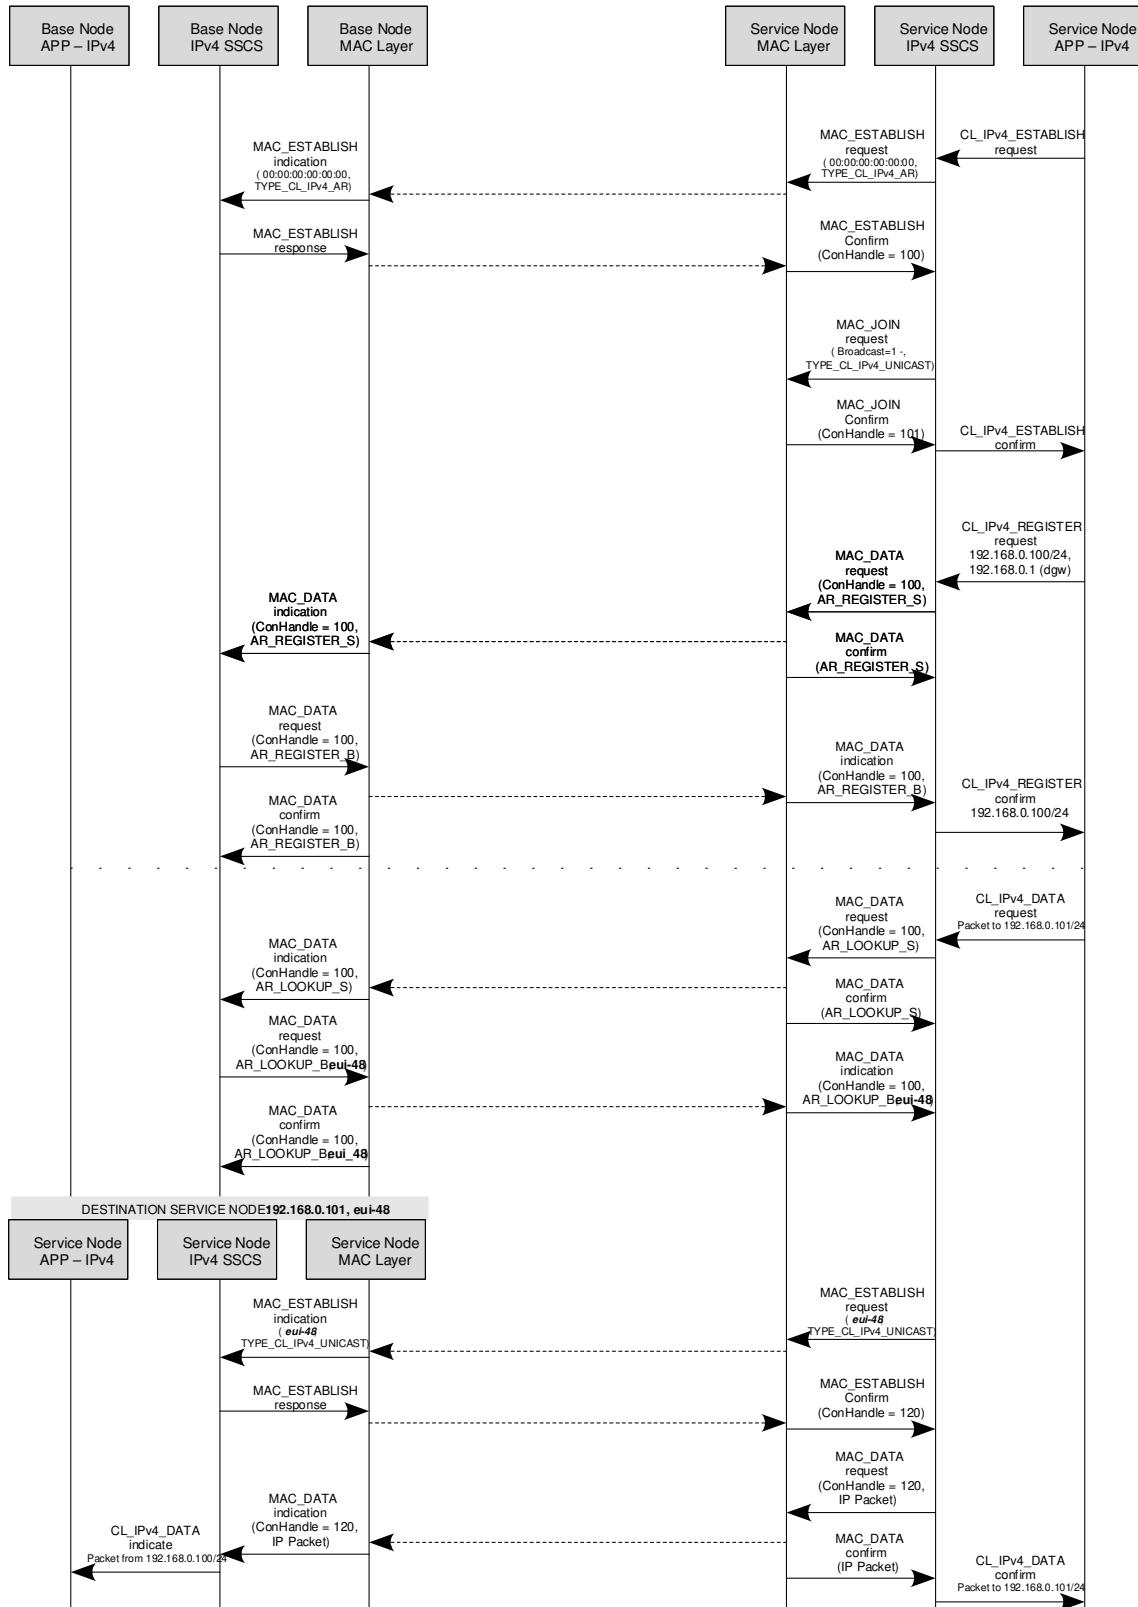
5050 b. To inform the service node MAC layer that IPv4 SSCS is ready to receive all IPv4 broadcasts packets.
5051 Note the difference between broadcast and multicast. To join a multicast group, the service node will
5052 need to inform the base node of the group it wants to join. This is illustrated in section A.2

5053 2. The IPv4_ layer, once the IPv4 SSCS is established, needs to register its IP address in the base node. To do
5054 so, it will use the already established connection.

5055 3. Whenever the IPv4_ needs to deliver an IPv4 packet to a new destination IP address, the following two
5056 steps are to be done (in this example, the destination IP address is 192.168.0.101).

5057 a. As the IPv4 destination address is new, the IPv4 SSCS needs to request the EUI-48 associated to that
5058 IPv4 address. To do so, a lookup request message is sent to the base node.

5059 b. Upon the reception of the EUI-48, a new connection (type = TYPE_CL_IPv4_UNICAST) is established
5060 so that all IP packets to be exchanged between 192.168.0.100 and 192.168.0.101 will use that
5061 connection.



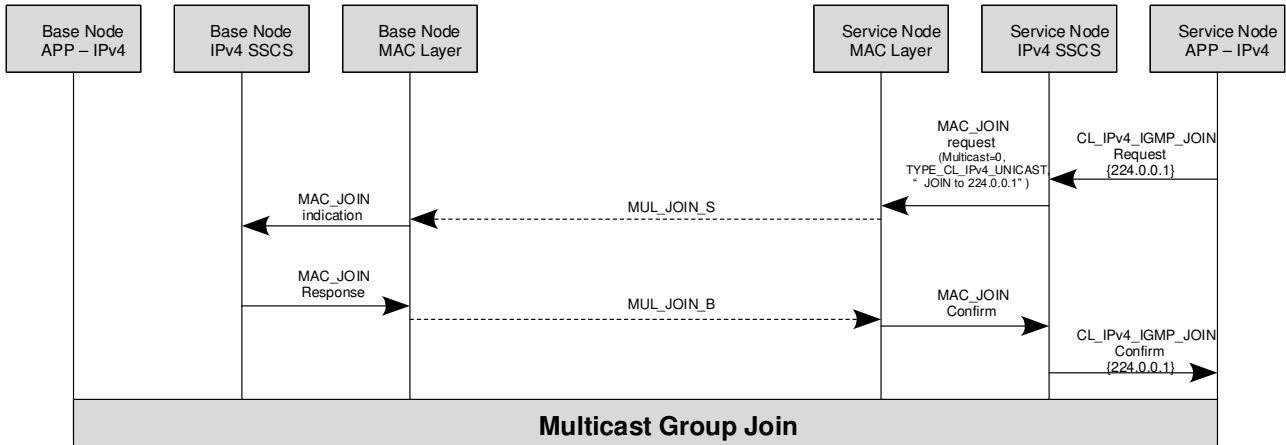
5062

5063

Figure H 1 - MSC of IPv4 SSCS services

5064 **H.2 Joining a multicast group**

5065 The figure below illustrates how a service node joins a multicast group. As mentioned before, main
 5066 difference between multicast and broadcast is related to the messages exchanged. For broadcast, the MAC
 5067 layer will immediately issue a MAC_JOIN.confirm primitive since it does not need to perform any end-to-
 5068 end operation. For multicast, the MAC_JOIN.confirm is only sent once the Control Packet transaction
 5069 between the service node and base node is complete.



5070

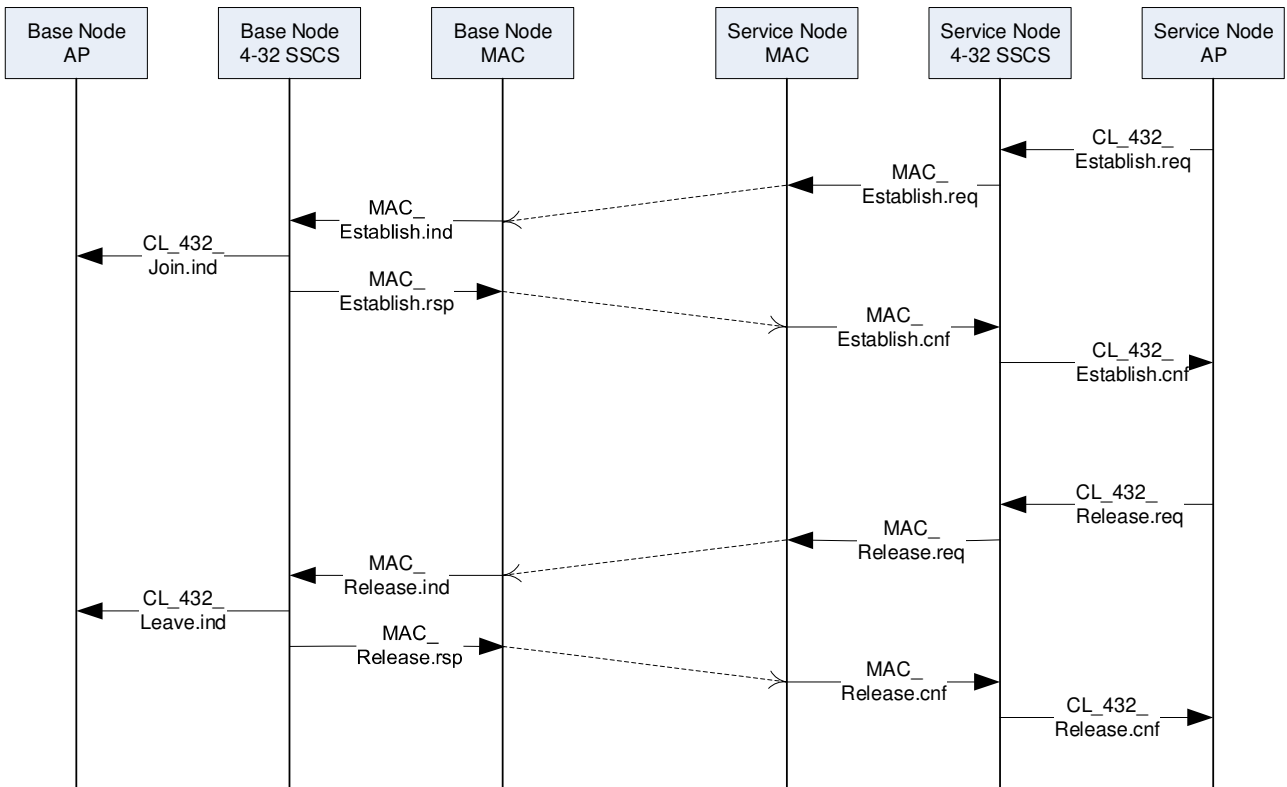
5071

5072

Figure H 2-Joining MS

5073 The MSC below shows the 432 connection establishment and release. 432 SSCS is connection oriented.
 5074 Before any 432_Data service can take place a connection establishment has to take place. The service node
 5075 upper layer request a connexion establishment to thez 432 SSCS by providing to it the device identifier as
 5076 parameter for the CL_432_Establish.request. With the help od the MAC layer services, the service node 432
 5077 SSCS request a connection establishment to the base node. This last one when the connection establishment
 5078 is successful, notifies to the upper layers that a service node has joined the network with the help of the
 5079 CL_432_Join.indication primitive and provides to the concerned service node a SSCS destination address in
 5080 addition to its own SSCS address with the help of the MAC_Establish.response which carries out these
 5081 parameters.

5082 The CL_432_release service ends the connection. It is requested by the service node upper layer to the 432
 5083 SSCS which perform it with the help of MAC layer primitives. At the base node side the 432 SSCS notifies the
 5084 end of the connection to the upper layer by a CL_432_Leave.indication.



5085

5086

5087

Figure H 3 - MSC 432 SCS services

5088
5089
5090

Annex I
(informative)
ARQ algorithm

5091 The algorithm described here is just a recommendation with good performance and aims to better describe
5092 how ARQ works. However manufacturers could use a different algorithm as long as it complies with the
5093 specification.

5094 When a packet is received the packet ID should be checked. If it is the expected ID and contains data, it
5095 shall be processed normally. If the packet does not contain data, it can be discarded. If the ID does not
5096 match with the one expected, it is from the future and fits in the input window, then for all the packets not
5097 received with ID from the last one received to this one, we can assume that they are lost. If the packet
5098 contains data, save that data to pass it to the CL once all the packets before have been received and
5099 processed by CL.

5100 If the packet ID does not fit in the input window, we can assume that it is a retransmission that has been
5101 delayed, and may be ignored.

5102 If there is any NACK all the packets with PKTID lower than the first NACK in the list have been correctly
5103 received, and they can be removed from the transmitting window. If there is not any NACK and there is an
5104 ACK, the packets before the received ACK have been received and can be removed from the transmission
5105 window. All the packets in the NACK list should be retransmitted as soon as possible.

5106 These are some situations for the transmitter to set the flush bit that may improve the average
5107 performance:

- 5108 • When the window of either the transmitter or the receiver is filled;
5109 • When explicitly requested by the CL;
5110 • After a period of time as a timeout.

5111 The receiver has no responsibility over the ACK send process other than sending them when the
5112 transmitter sets the flush bit. Although it has some control over the flow control by the window field. On
5113 the other hand the receiver is able to send an ACK if it improves the ARQ performance in a given scenario.
5114 One example of this, applicable in most cases, could be making the receiver send an ACK if a period of time
5115 has been passed since the last sent ACK, to improve the bandwidth usage (and omit the timeout flush in the
5116 transmitter). In those situations the transmitter still has the responsibility to interoperate with the simplest
5117 receiver (that does not send it by itself).

5118 It is recommended that the ARQ packet sender maintains a timer for every unacknowledged packet. If the
5119 packet cannot get successfully acknowledged when the timer expires, the packet will be retransmitted.
5120 This kind of timeout retry works independently with the NACK-initiated retries. After a pre-defined
5121 maximum number of timeout retries, it is strongly recommended to tear down the connection. This
5122 timeout and connection-teardown mechanism is to prevent the Node retry the ARQ packet forever. The
5123 exact number of the timeout values and the timeout retries are left for vendor's own choice.
5124

Annex J
(normative)
PHY backwards compatibility mechanism with PRIME v1.3.6


5125
5126
5127

5128 PRIME specification version V1.4 is an extension of version V1.3.6. The inclusion of new features, such as
5129 additional robust modes and a new frame type (Type B), implies that PRIME v1.4 compliant devices shall be
5130 able to support the following scenarios:

- 5131 1. Homogeneous networks which do not implement neither the new frame type (Type B) defined in
5132 Section 3.4 nor the additional robust modes (Robust DBPSK, Robust DQPSK).
- 5133 2. Homogeneous networks which implement the new frame type (Type B) defined in Section 3.4 as
5134 well as the additional robust modes (Robust DBPSK, Robust DQPSK).
- 5135 3. Mixed networks, composed of a combination of devices described in points (1) and (2) above.

5136 Cases (1) and (2) are trivial since the networks are homogeneous and all devices implement the same
5137 features. However, case (3) “Mixed networks” requires a specific mechanism that provides compatibility
5138 between PRIME compliant devices using different feature sets. Please note that compatibility between case
5139 (1) and case (2) devices could be trivially achieved forcing those devices with an extended set of features to
5140 ignore them and to use a more limited configuration (e.g., frame Type A and no robust modes).
5141 Nonetheless, the aim of this Annex is to define a backwards compatibility mechanism for mixed networks
5142 that allows devices with different feature sets to be part of the same network.

5143 Taking the PHY frame types into account, a backwards compatible frame (“BC frame”) is defined, as shown
5144 in Figure 138:



5145
5146

Figure 138 - Backwards Compatible PHY frame

5147 The BC frame is compatible with PRIME v1.3.6 frame at PHY level, since it is just a v1.3.6 PPDU
5148 (corresponding to a v1.4 Type A PPDU) encapsulating a v1.4 Type B PPDU.

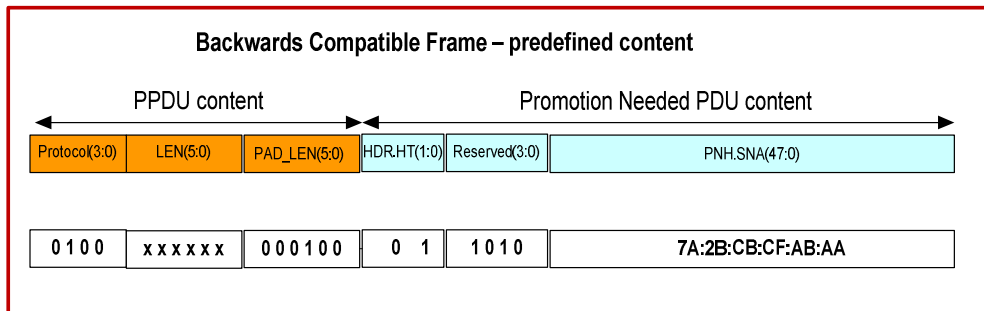
5149 BC frame predefined content is described below in

5150 Figure 139:

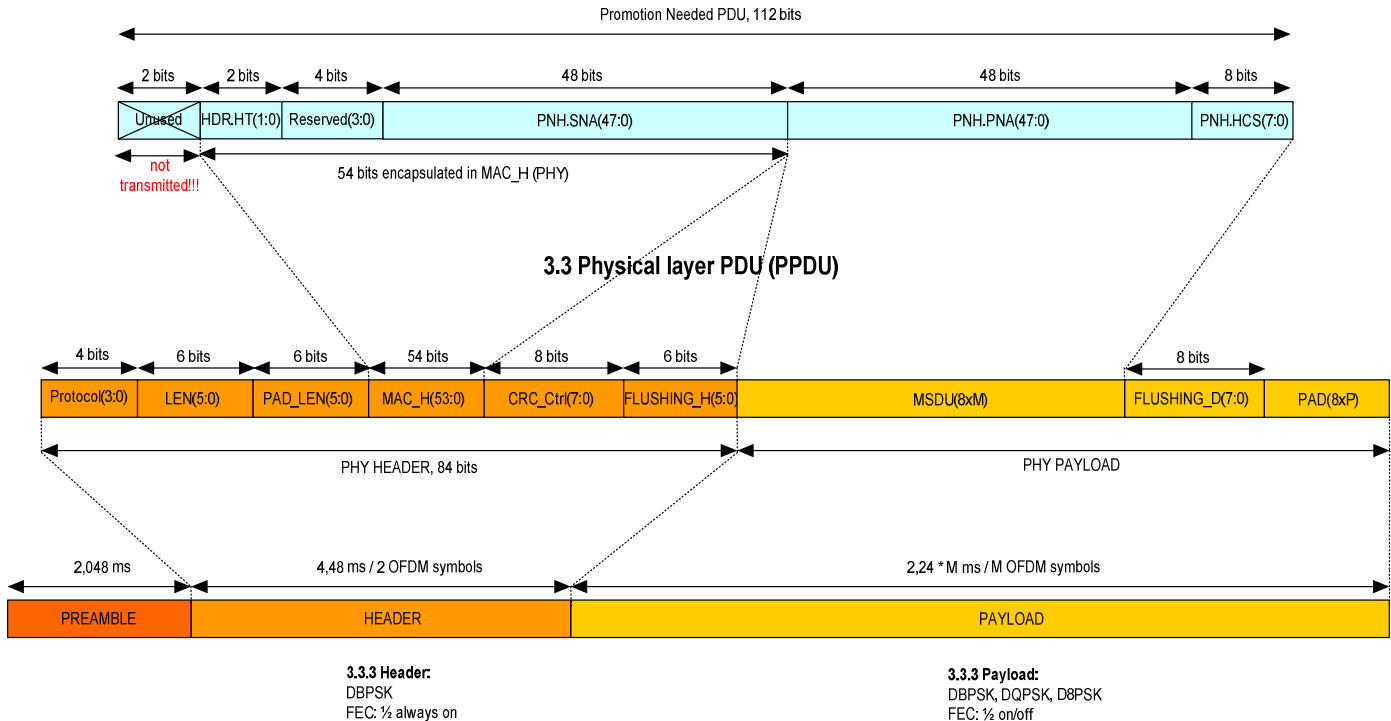
5151 a. PPDU content

- 5152 • Protocol [3:0]: Default transmission scheme, equal to DBPSK_CC

- 5153 • LEN[5:0]: Type A payload length (number of symbols in Type B header and Type B
- 5154 payload+4)
- 5155 b. PNPDU content
- 5156 • HDR.HT[1:0]: default PNPDU value, equal to “1”
- 5157 • Reserved[3:0]: predefined sequence, equal to “1010”
- 5158 • PNH.SNA[47:0]: predefined value, "7A:2B:CB:CF:AB:AA"



4.4.2 Promotion Needed PDU



5159

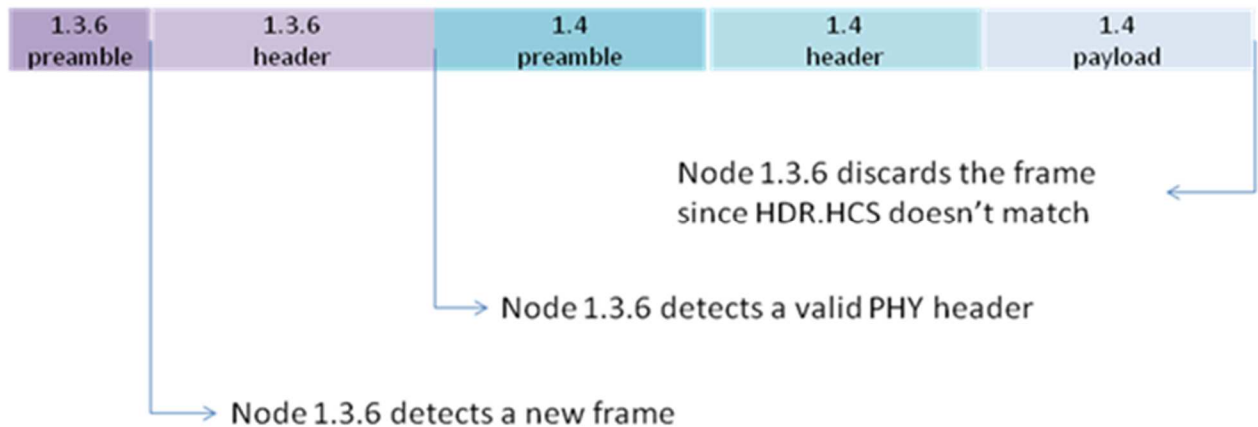
5160

Figure 139 - BC frame predefined content

5161 In mixed networks, the behavior of v1.4 and v1.3.6 devices upon reception of a BC frame will be different:

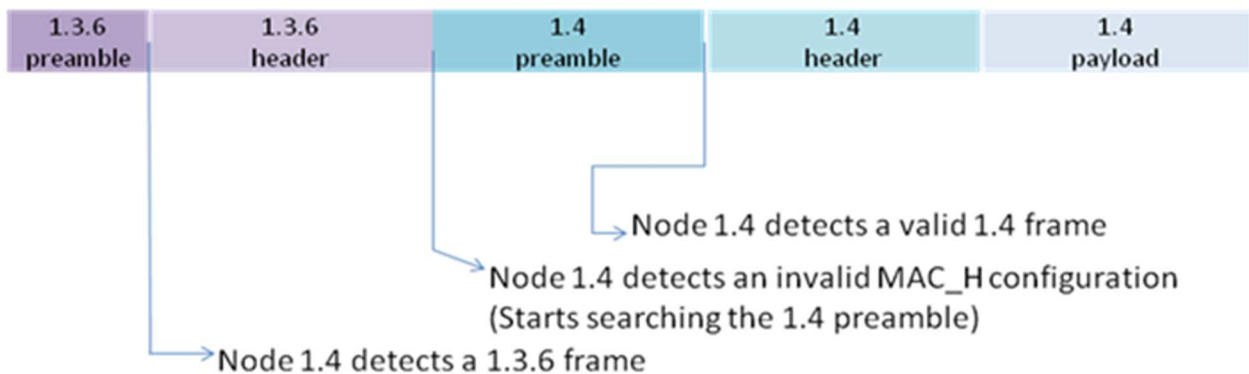
5162 v1.3.6 devices detect preamble and header. The content of the v1.3.6 header in a BC frame is a predefined
 5163 value (see

5164 1. Figure 139). The MAC of v1.3.6 devices will automatically discard the BC frame, but it will not
 5165 provoke any collisions while the frame is being transmitted (Figure 140).



5166
 5167 **Figure 140 - PHY BC frame detected by v1.3.6 devices**


5168 2. 1.4 devices will detect the v1.3.6 preamble and immediately after that the predefined MAC_H
 5169 configuration described above. Consequently, a v1.4 device will identify that a BC frame has been
 5170 received and shall start searching the v1.4 preamble and header (Figure 141):



5171
 5172 **Figure 141 - BC PHY frame detected by v1.4 devices**

5173 Two additional use cases have been added to this Annex for the sake of clarification:

5174 1. 1.3.6 frame received by a v1.4 node (Figure 142):



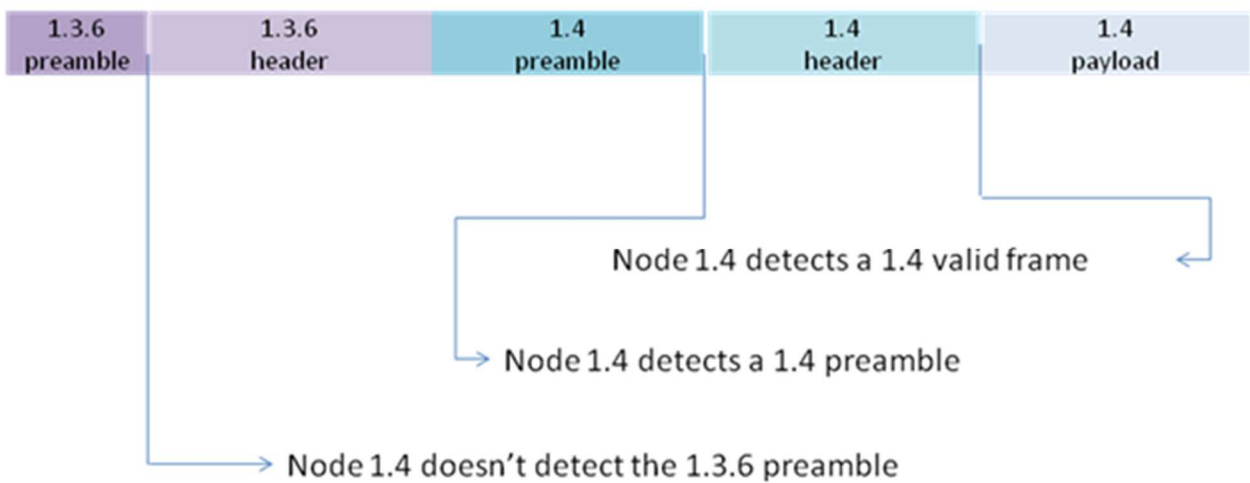
5175

5176

Figure 142 - v1.3.6 frame received by a v1.4 node

5177

5178 2. BC frame received by a v1.4 node in a very hard environment (Figure 143):



5179

5180

Figure 143 - BC PHY frame in noisy environment

5181

**Annex K
(normative)
MAC Backward Compatibility PDUs and Procedures**


5182
5183
5184

5185 **K.1 MAC PDU format**

5186 **K.1.1 Generic MAC PDU**

5187 In a network running in PRIME compatibility mode, all nodes shall use the standard Generic Mac header, as
5188 enumerated in Section 4.4.2.2, and the compatibility packet header (CPKT). The compatibility packet header
5189 is 6 bytes in length and its composition is shown in

5190 Figure 144. Table 151 enumerates the description of each field.



5191

5192

5193

Figure 144 - Compatibility Packet Header

Table 151 - Compatibility packet header fields

Name	Length	Description
<i>Reserved</i>	3 bits	Always 0 for this version of the specification. Reserved for future use.
CPKT.NAD	1 bit	No Aggregation at Destination <ul style="list-style-type: none"> • If CPKT.NAD=0 the packet may be aggregated with other packets at destination. • If CPKT.NAD=1 the packet may not be aggregated with other packets at destination.
CPKT.PRIO	2 bits	Indicates packet priority between 0 and 3.
CPKT.C	1 bits	Control <ul style="list-style-type: none"> • If CPKT.C=0 it is a data packet. • If CPKT.C=1 it is a control packet.

Name	Length	Description
CPKT.LCID / CPKT.CTYPE E	9 bits	Local Connection Identifier or Control Type <ul style="list-style-type: none"> • If CPKT.C=0, CPKT.LCID represents the Local Connection Identifier of data packet. • If CPKT.C=1, CPKT.CTYPE represents the type of the control packet.
CPKT.SID	8 bits	Switch identifier <ul style="list-style-type: none"> • If HDR.DO=0, CPKT.SID represents the SID of the packet source. • If HDR.DO=1, CPKT.SID represents the SID of the packet destination.
CPKT.LNID	14 bits	Local Node identifier. <ul style="list-style-type: none"> • If HDR.DO=0, CPKT.LNID represents the LNID of the packet source • If HDR.DO=1, CPKT.LNID represents the LNID of the packet destination.
CPKT.SPAD	1bit	Indicates if padding is inserted while encrypting payload. Note that this bit is only of relevance when Security Profile 1 (see 4.3.8.2.2) is used.
CPKT.LEN	9 bits	Length of the packet payload in bytes.

5194

5195 **K.1.1.1 MAC control packets**

5196 The CPKT.CTYPE field follows the same enumeration as the PKT.CTYPE field (see Table 19). Control packet
5197 retransmission shall follow the mechanisms described in Section 4.4.2.6.2.

5198 **K.1.1.1.1 Compatibility REG control packet (CREG, CPKT.CTYPE=1)**

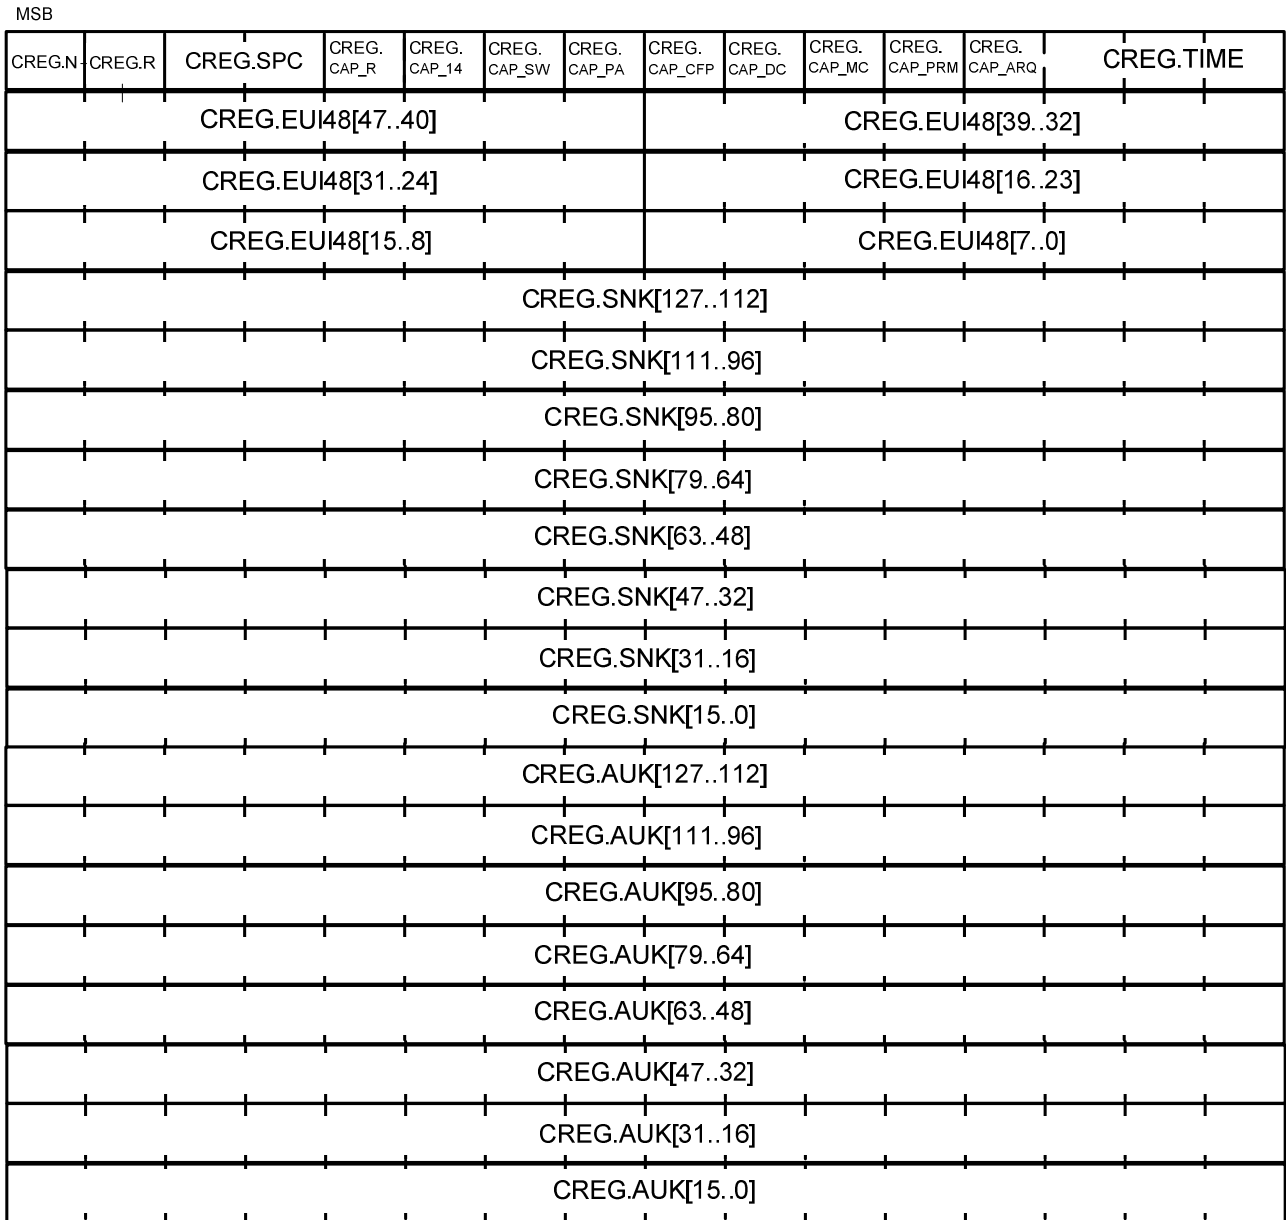
5199 The CREG control packet shall be used for registration requests (REG_REQ) in any case.

5200 REG and CREG control packets are distinguished based on the packet length. If the payload length is 8 or 40
5201 bytes, the payload is in CREG format; otherwise it is in REG format.

5202 The description of data fields of this control packet is described in Table 152 and Figure 145. The meaning of
5203 the packets differs depending on the direction of the packet. This packet interpretation is explained in Table
5204 153. These packets are used during the registration and unregistration processes in a compatibility mode
5205 network, as explained in Annex K.2.1 and K.2.2.

5206 The PKT.SID field is used in this control packet as the Switch where the Service Node is registering. The
5207 PKT.LNID field is used in this control packet as the Local Node Identifier being assigned to the Service Node
5208 during the registration process negotiation.

5209 The CREG.CAP_PA field is used to indicate the packet aggregation capability as discussed in Section 4.3.7. In
 5210 the uplink direction, this field is an indication from the registering Terminal Node about its own capabilities.
 5211 For the Downlink response, the Base Node evaluates whether or not all the devices in the cascaded chain
 5212 from itself to this Terminal Node have packet-aggregation capability. If they do, the Base Node shall set
 5213 CREG.CAP_PA=1; otherwise CREG.CAP_PA=0.



5214
 5215

Figure 145 - CREG control packet structure

5216

Table 152 - CREG control packet fields

Name	Length	Description
CREG.N	1 bit	Negative <ul style="list-style-type: none"> • CREG.N=1 for the negative register; • CREG.N=0 for the positive register. (see Table 153)
CREG.R	1 bit	Roaming <ul style="list-style-type: none"> • CREG.R=1 if Node already registered and wants to perform roaming to another Switch; • CREG.R=0 if Node not yet registered and wants to perform a clear registration process.
CREG.SPC	2 bits	Security Profile Capability for Data PDUs: <ul style="list-style-type: none"> • CREG.SPC=0 No encryption capability; • CREG.SPC=1 Security profile 1 capable device; • CREG.SPC=2 Security profile 2 capable device); • CREG.SPC=3 Security profile 3 capable device (not yet specified).
CREG.CAP_R	1 bit	Robust Mode <p>CREG REQ</p> <p>1 if the node is using a robust link to join the network;</p> <p>0 if the node is using a non-robust link to join the network</p> <p>CREG RSP</p> <p>the value is set to the same as CREG REQ frame</p> <p>CREG ACK</p> <p>the value is set to the same as CREG REQ frame</p>

Name	Length	Description
CREG.CAP_14	1 bit	<p>PRIME v1.4 Backward Compatibility Mode Capable</p> <p>1 (uplink) if the device is capable of using PRIME v1.4 backwards compatibility mode (i.e. this value is 1 for all PRIME v1.4 devices sending this message).</p> <p>1 (downlink) if the base node is acting in 1.4 backwards compatibility mode.</p> <p>0 if the device is a PRIME v1.3.6 device.</p>
CREG.CAP_SW	1 bit	<p>Switch Capable</p> <p>1 if the device is able to behave as a Switch Node;</p> <p>0 if the device is not.</p>
CREG.CAP_PA	1 bit	<p>Packet Aggregation Capability</p> <p>1 if the device has packet aggregation capability (uplink) if the data transit path to the device has packet aggregation capability (Downlink)</p> <p>0 otherwise.</p>
CREG.CAP_CFP	1 bit	<p>Contention Free Period Capability</p> <p>1 if the device is able to perform the negotiation of the CFP;</p> <p>0 if the device cannot use the Contention Free Period in a negotiated way.</p>
CREG.CAP_DC	1 bit	<p>Direct Connection Capability</p> <p>1 if the device is able to perform direct connections;</p> <p>0 if the device is not able to perform direct connections.</p>
CREG.CAP_MC	1 bit	<p>Multicast Capability</p> <p>1 if the device is able to use multicast for its own communications;</p> <p>0 if the device is not able to use multicast for its own communications.</p>

Name	Length	Description
CREG.CAP_PR M	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.
CREG.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.
CREG.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: CALV.TIME = 0 => 32 seconds; CALV.TIME = 1 => 64 seconds; CALV.TIME = 2 => 128 seconds ~ 2.1 minutes; CALV.TIME = 3 => 256 seconds ~ 4.2 minutes; CALV.TIME = 4 => 512 seconds ~ 8.5 minutes; CALV.TIME = 5 => 1024 seconds ~ 17.1 minutes; CALV.TIME = 6 => 2048 seconds ~ 34.1 minutes; CALV.TIME = 7 => 4096 seconds ~ 68.3 minutes.
CREG.EUI-48	48 bit	EUI-48 of the Node EUI-48 of the Node requesting the Registration.
CREG.SNK	128 bits	Encrypted Subnetwork key that shall be used to derive the Subnetwork working key
CREG.AUK	128 bits	Encrypted authentication key. This is a random sequence meant to act as authentication mechanism.

5217

Table 153 - CREG control packet types

Name	HDR.DO	CPKT.LNID	CREG.N	CREG.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.

Name	HDR.DO	CPKT.LNID	CREG.N	CREG.R	Description
REG_RSP	1	< 0x3FFF	0	R	Registration response. This packet assigns the CPCK.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

5218 Fields CREG.SNK and CREG.AUK are of significance only for REG_RSP and REG_ACK messages with Security
5219 Profile 1 (CREG.SCP=1). For all other message-exchange variants using the CREG control packet, these fields
5220 shall not be present reducing the length of payload.

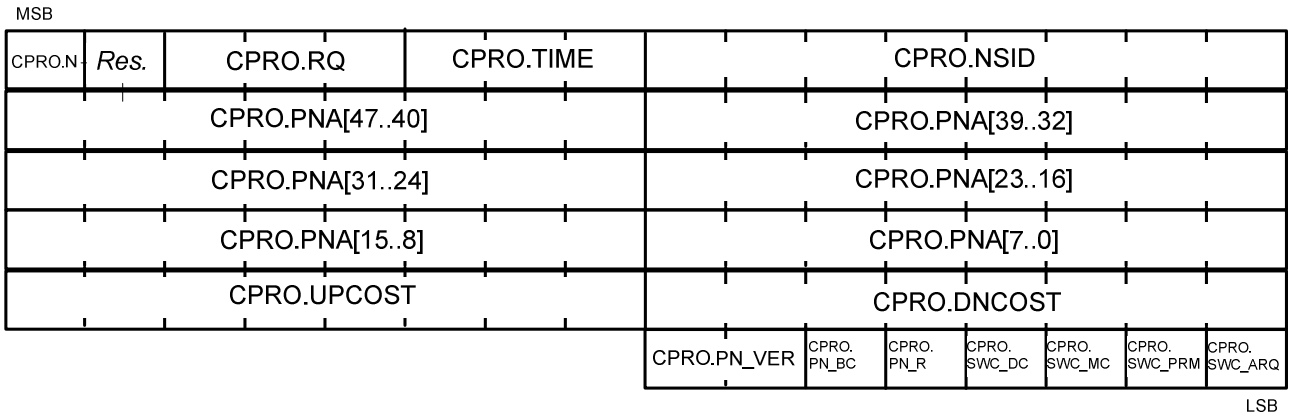
5221 In REG_RSP message, the CREG.SNK and CREG.AUK shall always be inserted encrypted with WK0.

5222 In the REG_ACK message, the CREG.SNK field shall be set to zero. The contents of the CREG.AUK field shall
5223 be derived by decrypting the received REG_RSP message with WK0 and re-encrypting the decrypted
5224 CREG.AUK field with SWK derived from the decrypted CREG.SNK and random sequence previously received
5225 in SEC control packets.

5226 **K.1.1.1.2 Compatibility PRO control packet (CPRO, CPKT.CTYPE = 3)**

5227 The compatibility promotion (CPRO) control packet is used by the base node and all service nodes to promote
5228 a Service Node from Terminal function to Switch function. The description of the fields of this packet is given
5229 in Table 154 and


5230 Figure 146. The meaning of the packet differs depending on the direction of the packet and on the values of
5231 the different types. Table 155 shows the different interpretation of the packets. The promotion process in
5232 backward compatibility mode is explained in more detail in Annex K.2.3 and K.2.3.1.



5233

5234

Figure 146 - CPRO_REQ_S control packet structure



5235

5236

Figure 147 - CPRO control packet structure

5237 Note that

5238 Figure 146 includes all fields as used by a CPRO_REQ_S message. All other messages are much smaller,
 5239 containing only CPRO.N, CPRO.RC, CPRO.TIME and CPRO.NSID as shown in Figure 147.

5240

Table 154 - CPRO control packet fields

Name	Length	Description
CPRO.N	1 bit	Negative CPRO.N=1 for the negative promotion CPRO.N=0 for the positive promotion
<i>Reserved</i>	1 bit	Reserved for future version of this protocol This shall be 0 for this version of the protocol.
CPRO.RQ	3 bits	Receive quality of the PNPDU message received from the Service Node requesting the Terminal to promote.
CPRO.TIME	3 bits	The ALV.TIME which is being used by the terminal which will become a switch. On a reception 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.

Name	Length	Description
CPRO.NSID	8 bits	<p>New Switch Identifier.</p> <p>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.</p>
CPRO.PNA	0 or 48 bits	<p>Promotion Need Address contains the EUJ-48 of the Terminal requesting the Service Node promotes to become a Switch.</p> <p>This field is only included in the PRO_REQ_S message.</p>
CPRO.UPCOST	0 or 8 bits	<p>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.</p> <p>This field is only included in the PRO_REQ_S message.</p>
CPRO.DNCOST	0 or 8 bits	<p>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.</p> <p>This field is only included in the PRO_REQ_S message.</p>
CPRO.PN_VER	2 bits	<p>Protocol version (PNH.VER) of the node represented by PRO.PNA.</p> <p>(This field is always zero for PRIME v1.3.6 nodes)</p>
CPRO.PN_BC	1 bit	<p>Backwards Compatibility mode of the node represented by PRO.PNA.</p> <p>1 if the device is backwards compatible with 1.3.6 PRIME 0 if it is not.</p> <p>(This field is always zero for PRIME v1.3.6 nodes)</p>
CPRO.PN_R	1 bit	<p>Robust mode compatibility of the node represented by PRO.PNA.</p> <p>1 if the device supports robust mode 0 if it is not</p> <p>(This field is always zero for PRIME v1.3.6 nodes)</p>

Name	Length	Description
CPRO.SWC_DC	1 bit	Direct Connection Switching Capability 1 if the device is able to behave as Direct Switch in direct connections. 0 otherwise
CPRO.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
CPRO.SWC_PR M	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.
CPRO.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.

5241

Table 155 - CPRO control packet types


Name	HDR.DO	CPRO.N	CPRO.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 CPRO.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.

Name	HDR.DO	CPRO.N	CPRO.NSID	Description
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.

5242

5243 **K.1.1.1.3 Compatibility BSI control packet (CBSI, CPKT.CTYPE = 4)**

5244 The Compatibility Beacon Slot Information (CBSI) control packet is only used by the Base Node and Switch
 5245 Nodes. It is used to exchange information that is further used by a Switch Node to transmit its beacon. The
 5246 description of the fields of this packet is given in Table 156 and Figure 148. The meaning of the packet differs
 5247 depending on the direction of the packet and on the values of the different types. Table 157 represents the
 5248 different interpretation of the packets. The promotion process is explained in more detail in 4.6.3.



5249

5250 **Figure 148 - CBSI control packet structure**

5251 **Table 156 - CBSI control packet fields**

Name	Length	Description
Reserved	5 bits	Reserved for future version of this protocol. In this version, this field should be initialized to 0.
CBSI.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 FRQ = 7 => Reserved


Name	Length	Description
CBSI.SLT	3 bits	Beacon Slot to be used by target Switch 0 – 4: non-robust mode beacon slot 5 – 6: robust mode beacon slot
CBSI.SEQ	5 bits	The Beacon Sequence number when the specified change takes effect.

5252 **Table 157 - CBSI control message types**

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

5253
5254 **K.1.1.1.4 Compatibility FRA control packet (CFRA, CPKT.CTYPE = 5)**

5255 This control packet is broadcast from the Base Node and relayed by all Switch Nodes to the entire
5256 Subnetwork. It is used by switches transmitting CBCN compatibility beacons, and the terminal nodes directly
5257 attached to them, to learn about upcoming frame changes. The description of fields of this packet is given in
5258 Table 158 and Figure 149. Table 159 shows the different interpretations of the packets.



5259
5260 **Figure 149 - CFRA control packet structure**

5261 **Table 158 - CFRA control packet fields**

Name	Length	Description
CFRA.TYP	2 bits	0: Beacon count change 1: CFP duration change 2: Reserved for PRIME v1.4 FRA message (see Section 4.4.2.6.6)
Reserved	4 bits	Reserved for future version of this protocol. In this version, this field should be initialized to 0.
CFRA.CFP	10 bits	Offset of CFP from start of frame
CFRA.SEQ	5 bits	The Beacon Sequence number when the specified change takes effect.
CFRA.BCN	3 bits	Number of beacons in a frame


5262

Table 159 - CFRA control packet types

Name	CFRA.TYP	Description
FRA_BCN_IND	0	Indicates changes to frame structure due to change in 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.

5263 **K.1.1.1.5 Compatibility ALV control packet (CALV, CPKT.CTYPE = 7)**

5264 In a compatibility mode network, the CALV control message is used exclusively for Keep-Alive signaling
 5265 between a Service Node, the Service Nodes above it and the Base Node. The message exchange is
 5266 bidirectional, that is, a message is periodically exchanged in each direction. The structure of these messages
 5267 is shown in Figure 150 and Table 160. The different Keep-Alive message types are shown in Table 161. The
 5268 compatibility keep-alive process is shown in Annex K.2.5.



5269

5270

5271

Figure 150 - CALV Control packet structure

Table 160 - CALV control message fields

Name	Length	Description
CALV.RXCNT	3 bits	Modulo 8 counter to indicate number of received CALV messages.
CALV.TXCNT	3 bits	Modulo 8 counter to indicate number of transmitted CALV messages.
<i>Reserved</i>	7 bits	Should always be encoded as 0 in this version of the specification.
CALV.TIME	3 bits	Time to wait for an ALV_B messages before assuming the Service Node has been unregistered by the Base Node. CALV.TIME = 0 => 32 seconds; CALV.TIME = 1 => 64 seconds; CALV.TIME = 2 => 128 seconds ~ 2.1 minutes; CALV.TIME = 3 => 256 seconds ~ 4.2 minutes; CALV.TIME = 4 => 512 seconds ~ 8.5 minutes; CALV.TIME = 5 => 1024 seconds ~ 17.1 minutes; CALV.TIME = 6 => 2048 seconds ~ 34.1 minutes; CALV.TIME = 7 => 4096 seconds ~ 68.3 minutes.
CALV.SSID	8 bits	For a Terminal, this should be 0xFF. For a Switch, this is its Switch Identifier.

5272

Table 161 - Keep-Alive control packet types

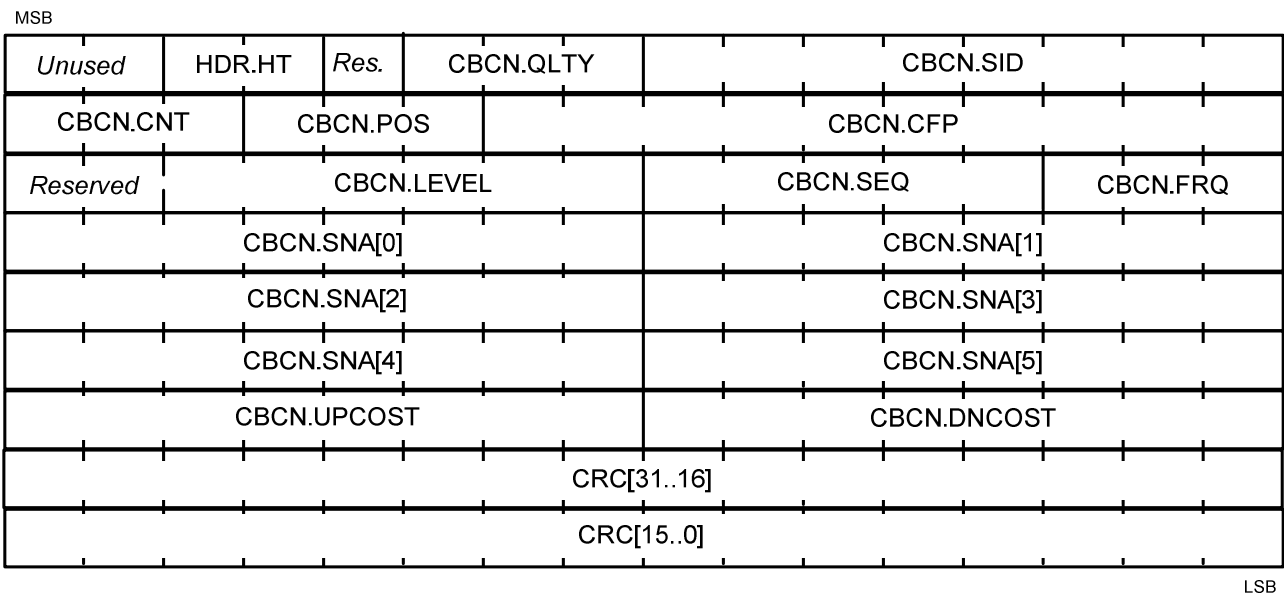
Name	HDR.DO	Description
ALV_S	0	Keep-Alive message from a Service Node
ALV_B	1	Keep-Alive message from the Base Node

5273

5274 **K.1.2 Compatibility Beacon PDU (CBCN)**

5275 In a compatibility mode network, the compatibility beacon PDU (CBCN) is transmitted by the base node and
 5276 some of the Switch devices on the Subnetwork (see table in section 4.9.3.2).

5277 Figure 151 below shows contents of a CBCN beacon.



5278

LSB

5279

Figure 151 - Beacon PDU structure

5280 Table 162 shows the CBCN PDU fields.

5281

Table 162 - Beacon PDU fields

Name	Length	Description
Unused	2 bits	Unused bits which are always 0; included for alignment with MAC_H field in PDU header (Fig 7, Section 3.3.3).
HDR.HT	2 bits	Header Type HDR.HT = 2 for Beacon PDU
Reserved	1 bit	Always 0 for this version of the specification. Reserved for future use.

Name	Length	Description
CBCN.QLTY	3 bits	Quality of round-trip connectivity from this Switch Node to the Base Node. CBCN.QLTY=7 for best quality (Base Node or very good Switch Node), CBCN.QLTY=0 for worst quality (Switch having unstable connection to Subnetwork)
CBCN.SID	8 bits	Switch identifier of transmitting Switch
CBCN.CNT	3 bits	Number of beacon-slots in this frame
CBCN.SLT	3 bits	Beacon-slot in which this BPDU is transmitted CBCN.SLT=0 is reserved for the Base Node
CBCN.CFP	10 bits	Offset of CFP from start of frame CBCN.CFP=0 indicates absence of CFP in a frame. (CBCN. CFP includes robust beacon slots)
Reserved	1 bit	Always 0 for this version of the specification. Reserved for future use.
CBCN.LEVEL	6 bits	Hierarchy of transmitting Switch in Subnetwork
CBCN.SEQ	5 bits	Sequence number of this BPDU in super frame. Incremented for every beacon the Base Node sends and is propagated by Switch through its BPDU such that entire Subnetwork has the same notion of sequence number at a given time.
CBCN.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 7 = Reserved
CBCN.SNA	48 bits	Subnetwork identifier in which the Switch transmitting this BPDU is located

Name	Length	Description
CBCN.UPCOST	8 bits	<p>Total uplink cost from the transmitting Switch Node to the Base Node. The cost of a single hop is calculated based on modulation scheme used on that hop in uplink direction. Values are derived as follows:</p> <p style="margin-left: 40px;">8PSK = 0 QPSK = 1 BPSK = 2 8PSK_F = 1 QPSK_F = 2 BPSK_F = 4</p> <p>The Base Node will transmit in its beacon a CBCN.UPCOST of 0. A Switch Node will transmit in its beacon the value of CBCN.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 CBCN.UPCOST the maximum value of CBCN.UPCOST should be used.</p>
CBCN.DNCOST	8 bits	<p>Total Downlink cost from the Base Node to the transmitting Switch Node. The cost of a single hop is calculated based on modulation scheme used on that hop in Downlink direction. Values are derived as follows:</p> <p style="margin-left: 40px;">8PSK 0 QPSK 1 BPSK 2 8PSK_F 1 QPSK_F 2 BPSK_F 4</p> <p>The Base Node will transmit in its beacon a CBCN.DNCOST of 0. A Switch Node will transmit in its beacon the value of CBCN.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 CBCN.DNCOST the maximum value of CBCN.DNCOST should be used.</p>
CRC	32 bits	<p>The CRC shall be calculated with the same algorithm as the one defined for the CRC field of the MAC PDU (see section 0 for details). This CRC shall be calculated over the complete BPDU except for the CRC field itself.</p>

5283 The CBCN BPDU is also used to detect when the uplink Switch is no longer available either by a change in the
 5284 characteristics of the medium or because of failure etc. The rules in section 4.4.4 apply.

5285 **K.2 MAC procedures**

5286 **K.2.1 Registration process**

5287 The initial Service Node start-up (4.3.1) is followed by a Registration process. A Service Node in a
 5288 *Disconnected* functional state shall transmit a registration control packet to the Base Node in order to get
 5289 itself included in the Subnetwork.

5290 All beacons shall be sent in CBCN format.

5291 Since no LNID or SID is allocated to a Service Node at this stage, the CPKT.LNID field shall be set to all 1s and
 5292 the CPKT.SID field shall contain the SID of the Switch Node through which it seeks attachment to the
 5293 Subnetwork.

5294 Base Nodes may use a Registration request as an authentication mechanism. However this specification does
 5295 not recommend or forbid any specific authentication mechanism and leaves this choice to implementations.

5296 For all successfully accepted Registration requests, the Base Node shall allocate an LNID that is unique within
 5297 the domain of the Switch Node through which the attachment is realized. This LNID shall be indicated in the
 5298 PKT.LNID field of response (REG_RSP). The assigned LNID, in combination with the SID of the Switch Node
 5299 through which the Service Node is registered, would form the NID of the registering Node.

5300 Based on the flag CBNC.CAP_14 a service node knows whether it is registered to a PRIME v1.4 (standard or
 5301 compatibility mode) or a PRIME v1.3.6 network. Registration is a three-way process. The REG_RSP shall be
 5302 acknowledged by the receiving Service Node with a REG_ACK message. The same format is used for the
 5303 REG_RSP as for the REG_REQ.

5304 Figure 152 represents a successful Registration process and Figure 153 shows a Registration request that is
 5305 rejected by the Base Node. Details on specific fields that distinguish one Registration message from the other
 5306 are given in Table 21 and Table 153.

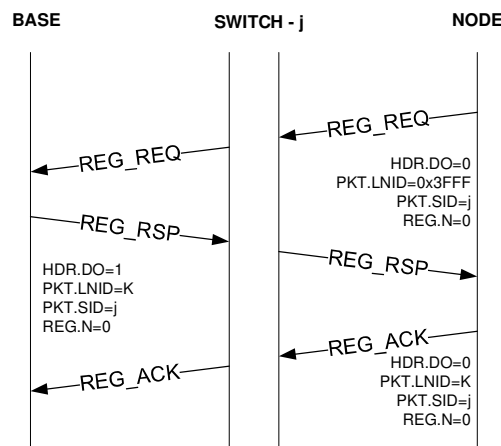


Figure 152 - Registration process accepted

5307

5308




Figure 153 - Registration process rejected

5309

5310

5311 When assigning an LNID, the Base Node shall not reuse an LNID released by an unregister process until after
 5312 (*macCtrlMsgFailTime* + *macMinCtlReTxTimer*) seconds, to ensure that all retransmit packets have left the
 5313 Subnetwork. Similarly, the Base Node shall not reuse an LNID freed by the Keep-Alive process until T_{keep_alive}
 5314 seconds have passed, using the last known acknowledged T_{keep_alive} value, or if larger, the last unacknowledged
 5315 T_{keep_alive} , for the Service Node using the LNID.

5316 During network startup where the whole network is powered on at once, there will be considerable
 5317 contention for the medium. It is recommended, but optional, that randomness is added to the first
 5318 transmission of REQ_REQ and all subsequent retransmissions. A random delay of maximum duration of 10%
 5319 of *macMinCtlReTxTimer* may be imposed before the first REG_REQ message, and a similar random delay of
 5320 up to 10% of *macMinCtlReTxTimer* may be added to each retransmission.

5321 **K.2.2 Unregistering process**

5322 The unregistering process follows the description in Section 4.6.2. All nodes use compatibility mode
 5323 unregistration packets (CREG).

5324 **K.2.3 Promotion process**

5325 A Node that cannot reach any existing Switch may send promotion-needed frames so that a Terminal can be
 5326 promoted and begin to switch. During this process, a Node that cannot reach any existing Switch may send
 5327 PNPDU's so that a nearby Terminal can be promoted and begin to act as a Switch. During this process, a
 5328 Terminal will receive PNPDU's and at its discretion, generate compatibility mode PRO_REQ control packets to
 5329 the Base Node. In a compatibility mode network no standard PRO messages are used but only CPRO and CBSI
 5330 messages.

5331 The Base Node examines the promotion requests during a period of time. It may use the address of the new
 5332 Terminal, provided in the promotion-request packet, to decide whether or not to accept the promotion. It
 5333 will decide which Node shall be promoted, if any, sending a promotion response. The other Nodes will not
 5334 receive any answer to the promotion request to avoid Subnetwork saturation. Eventually, the Base Node
 5335 may send a rejection if any special situation occurs. If the Subnetwork is specially preconfigured, the Base
 5336 Node may send Terminal Node promotion requests directly to a Terminal Node.

5337 When a Terminal Node requests promotion, the CPRO.NSID field in the PRO_REQ_S message shall be set to
 5338 all 1s. The PRO.NSID field shall contain an LSID allocated to the promoted Node in the PRO_REQ_B message.

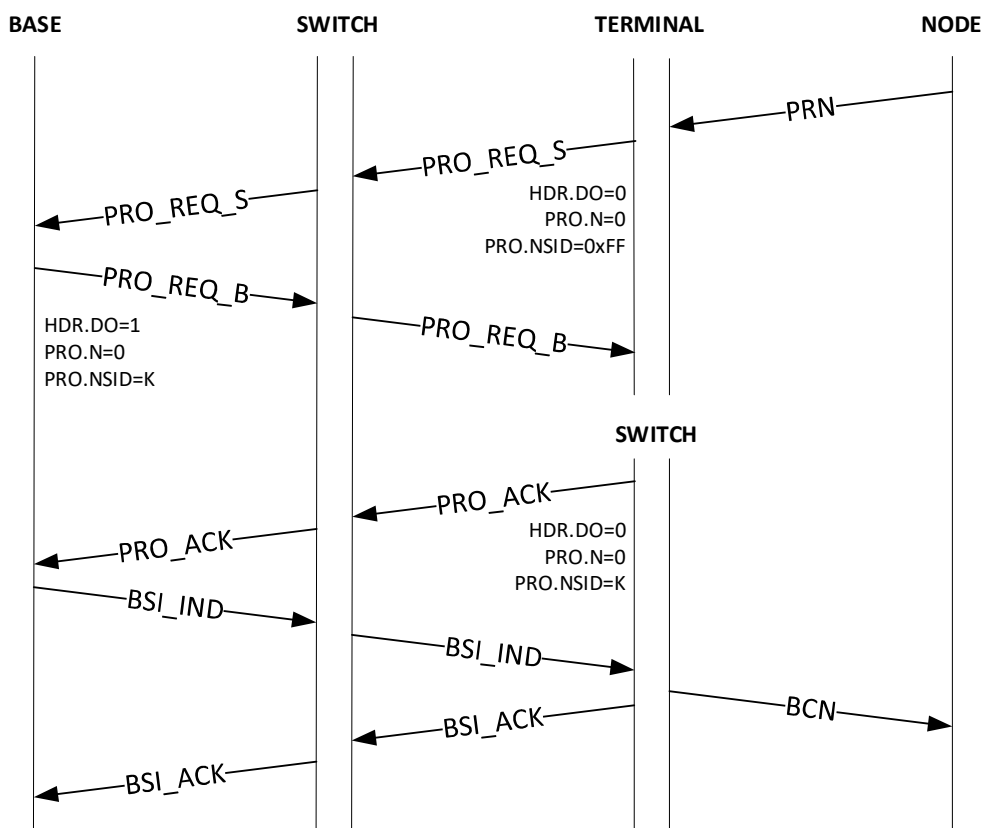
5339 The acknowledging Switch Node shall set the CPRO.NSID field in its PRO_ACK to the newly allocated LSID.
 5340 This final PRO_ACK shall be used by intermediate Switch Nodes to update their switching tables as described
 5341 in 4.3.5.2.

5342 When reusing LSIDs that have been released by a demotion process, the Base Node should not allocate the
 5343 LSID until after $(macCtrlMsgFailTime + macMinCtrlReTxTimer)$ seconds to ensure all retransmit packets that
 5344 might use that LSID have left the Subnetwork. Similarly, the Base Node shall not reuse an LNID freed by the
 5345 Keep-Alive process until T_{keep_alive} seconds have passed, using the last known acknowledged T_{keep_alive} value,
 5346 or if larger, the last unacknowledged T_{keep_alive} , for the Service Node using the LNID.

5347 After the base node receives the PRO_ACK, the Base Node sends a BSI_IND to the service node. The encoding
 5348 of the Beacon is decided using the beacon slot, if the beacon slot is 5 or 6, the encoding shall be DBPSK_R.
 5349 The service node shall respond with the corresponding BSI_ACK.

5350 The base node can use BSI_IND with two purposes:

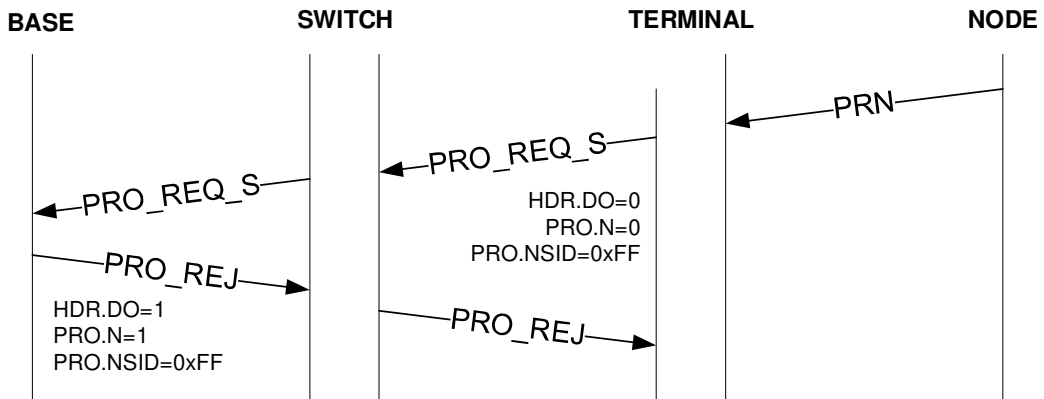
- 5351 • Change the allocation of the transmitted beacon. Only if the robustness of the beacon does not
 5352 change.
- 5353 • Start double switching by sending a second beacon in the other modulation.
- 5354 • After a switch is double switching the next BSI_IND shall change the transmission properties of the
 5355 robust beacon if slot is 5 or 6 and of the non-robust beacon otherwise.



5356

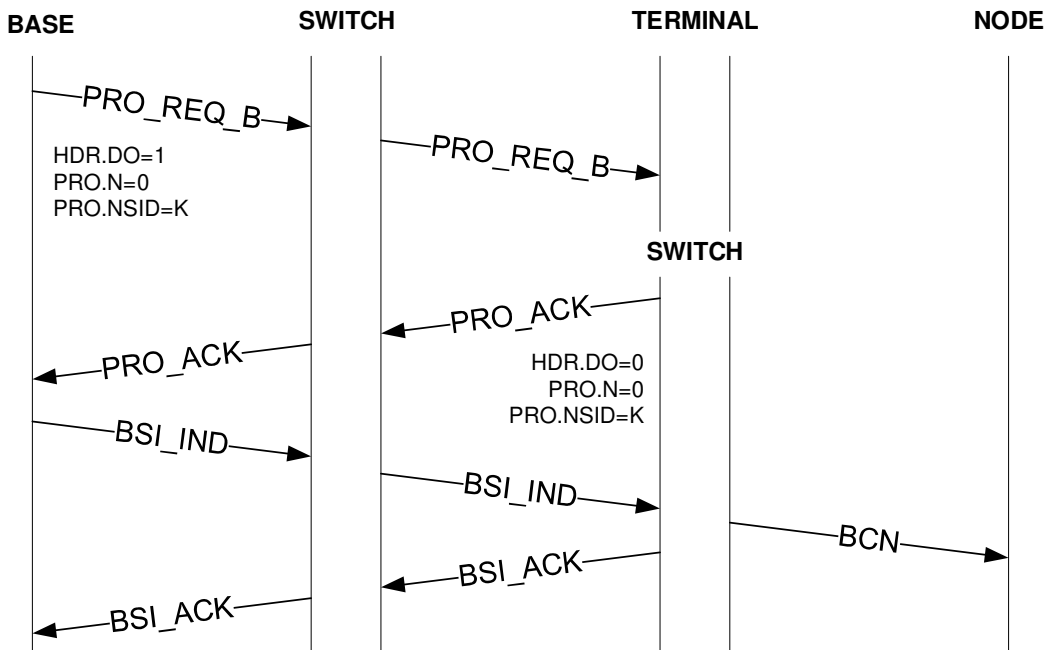
5357

Figure 154 - Promotion process initiated by a Service Node



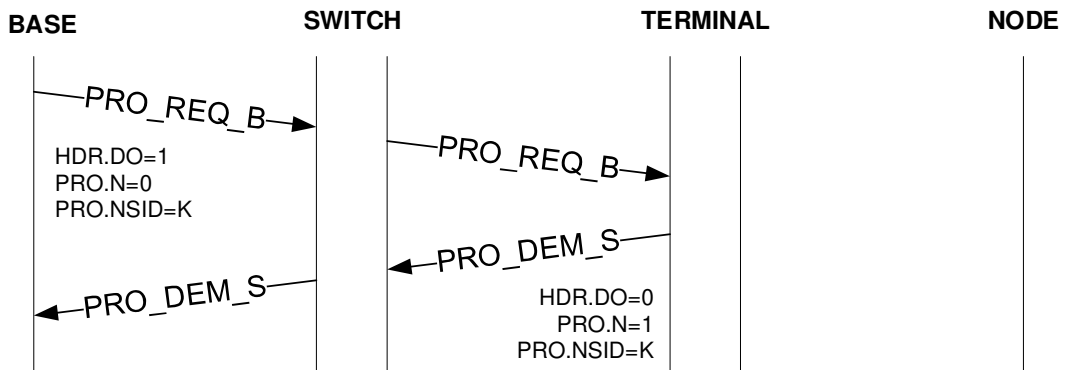
5358
5359

Figure 155 - Promotion process rejected by the Base Node



5360
5361

Figure 156 - Promotion process initiated by the Base Node



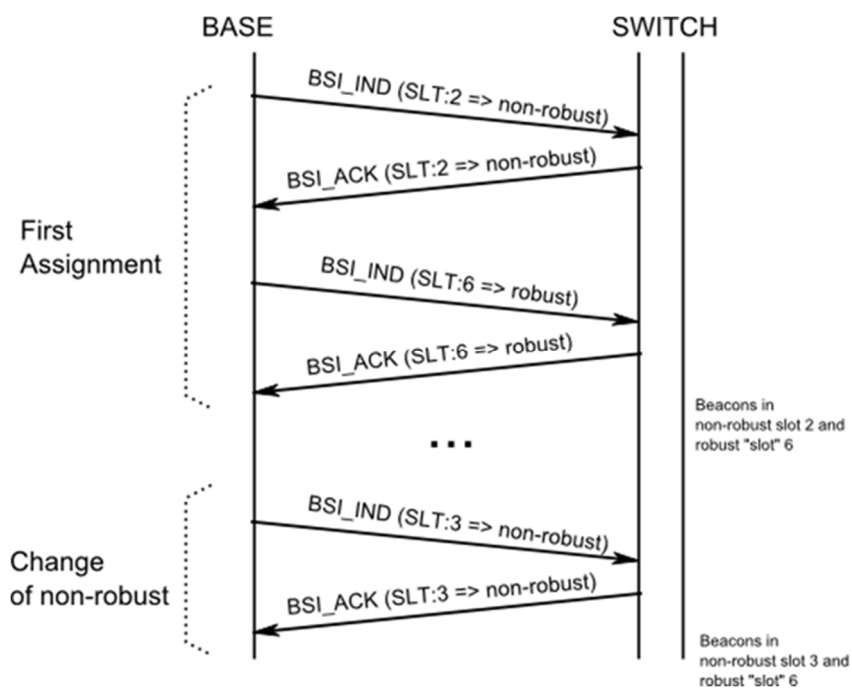
5362
5363

Figure 157 - Promotion process rejected by a Service Node

5364 **K.2.3.1 Double switching**

5365 Every time a Base Node promotes a node to act as robust switch, it shall start two BSI procedures to promote
 5366 the Service Node so it has two beacon slots assigned, one robust and one non-robust.

5367 One of the BSI_IND shall have a beacon slot in the range 0-4 for the non-robust beacon (DBPSK_CC) and the
 5368 other one shall send the beacon slot in the range 5-6 for the robust beacon (DBPSK_R). For future changes
 5369 of the BSI information the rule to separate the robust and non-robust beacons shall be the range of the
 5370 beacon slot



5371
 5372 **Figure 158 - Double switching BSI message exchange**


5373 **K.2.4 Demotion process**

5374 The Base Node or a Switch Node may decide to discontinue a switching function at anytime. The demotion
 5375 process provides for such a mechanism. In a compatibility mode network, only CPRO control packets are used
 5376 for all demotion transactions.

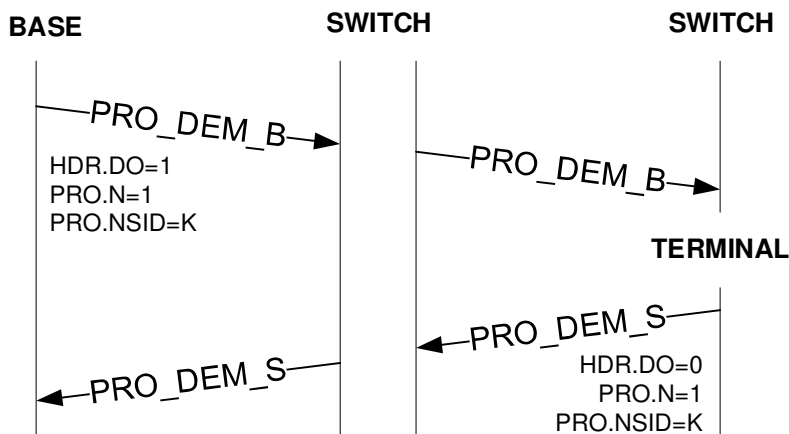
5377 The CPRO.NSID field shall contain the SID of the Switch Node that is being demoted as part of the demotion
 5378 transaction. The PRO.PNA field is not used in any demotion process transaction and its contents are not
 5379 interpreted at either end.

5380 Following the successful completion of a demotion process, a Switch Node shall immediately stop the
 5381 transmission of beacons and change from a *Switch* functional state to a *Terminal* functional state. The Base
 5382 Node may reallocate the LSID and Beacon Slot used by the demoted Switch after (*macCtrlMsgFailTime* +
 5383 *macMinCtReTxTimer*) seconds to other Terminal Nodes requesting promotion.

5384 The present version of this specification does not specify any explicit message to reject a demotion requested
 5385 by a peer at the other end.



5386
 5387 **Figure 159 - Demotion process initiated by a Service Node**



5388
 5389 **Figure 160 - Demotion process initiated by the Base Node**

5390 **K.2.5 Keep-Alive process**

5391 The Keep-Alive process in a compatibility mode network is fundamentally different from the Keep-Alive
 5392 process used in a standard PRIME v1.4 network. It is based on the PRIME v1.3.6 end-to-end process. The
 5393 Keep-Alive process is used to detect when a Service Node has left the Subnetwork because of changes to the
 5394 network configuration or because of fatal errors it cannot recover from.

5395 When the Service Node receives the REG_RSP packet it uses the REG.TIME/CREG.TIME field to start a timer
 5396 T_{keep_alive} . For every ALV_B it receives, it restarts this timer using the value from CALV.TIME. The encoding of
 5397 CALV.TIME is specified in Table 160. It should also send an ALV_S to the Base Node. If the timer ever expires,
 5398 the Service Node assumes it has been unregistered by the Base Node. The message PRO_REQ does also reset
 5399 the Keep-Alive timer to the CPRO.TIME value.

5400 Each switch along the path of a ALV_B message takes should keep a copy of the CPRO.TIME and then
 5401 CALV.TIME for each Switch Node below it in the tree. When the switch does not receive an ALV_S message

5402 from a Service Node below it for $T_{\text{keep_alive}}$ as defined in CPRO.TIME and CALV.TIME it should remove the
5403 Switch Node entry from its switch table. See section 4.3.5.2 for more information on the switching table.
5404 Additionally a Switch Node may use the REG.TIME/CREG.TIME and CALV.TIME to consider also every Service
5405 Node Registration status and take it into account for the switching table.

5406 For every ALV_S or ALV_B message sent by the Base Node or Service Node, the counter CALV.TXCNT should
5407 be incremented before the message is sent. This counter is expected to wrap around. For every ALV_B or
5408 ALV_S message received by the Service Node or the Base Node the counter CALV.RXCNT should be
5409 incremented. This counter is also expected to wrap around. These two counters are placed into the ALV_S
5410 and ALV_B messages. The Base Node should keep a CALV.TXCNT and CALV.RXCNT separated counter for each
5411 Service Node. These counters are reset to zero in the Registration process.

5412 The algorithm used by the Base Node to determine when to send ALV_B messages to registered Service
5413 Nodes and how to determine the value CALV.TIME and REG.TIME/CREG.TIME is not specified here.

5414 **K.2.6 Connection management**

5415 The processes follow the standard processes described in section 4.3

5416 **K.2.7 Multicast group management**

5417 The processes follow the standard processes described in section 4.6.7. The base node shall not send any
5418 MUL_SW_LEAVE_B to PRIME v1.3.6 service nodes, as the PRIME v1.3.6 switches implement a different
5419 mechanism for multicast group tracking.

5420 **K.2.8 Robustness Management**

5421 Robustness management is not performed between devices running legacy version of protocol.

5422 **K.2.9 Channel allocation and deallocation**

5423 The process follows the description in Section 4.6.9. The Base Node shall send a CFRA broadcast packet.

**Annex L
(Informative)
Type A, Type B PHY frames and Robust modes**

5424
5425
5426

5427 The following is a recommendation about how to combine the two PHY frame formats defined by PRIME
5428 with the available payload transmission schemes. As a general guideline, preamble and header shall be at
5429 least as robust as the payload.

5430 Type A and Type B PRIME PHY frames specify different Preamble lengths and Header formats:

- 5431 - TYPE A PHY frames, as described in Figure 3, comprise a "Preamble A" lasting 2.048 ms and a "Header
5432 A" with a length equal to two OFDM symbols (2 x 2.24 ms).
- 5433 - TYPE B PHY frames, as described in Figure 4, **achieve higher robustness** by means of a "Preamble B"
5434 lasting 8.192 ms and a "Header B" with a length equal to four OFDM symbols (4 x 2.24 ms)

5435 Table 163 shows all possible combinations, recommendations [OK / NOK] are based on the fact that preamble
5436 and header shall be at least as robust as the payload.

5437 **Table 163 - PHY frame types and Payload transmission schemes**

HEADER and PREAMBLE	PAYLOAD							
	Robust DBPSK	Robust DQPSK	DBPSK_CC	DBPSK	DQPSK_CC	DQPSK	D8PSK_CC	D8PSK
Type A (short preamble, short header)	NOK	NOK	OK	OK	OK	OK	OK	OK
Type B (long preamble, long header)	OK	OK	OK	OK	OK	OK	OK	OK

5438

5439 List of authors (by alphabetical order)

- 5440 *Ankou, Auguste (Itron)*
- 5441 *Arzuaga, Aitor (ZIV)*
- 5442 *Arzuaga, Txetxu (ZIV)*
- 5443 *Berganza, Inigo (Iberdrola)*
- 5444 *Bertoni, Guido (STMicroelectronics)*
- 5445 *Bisaglia, Paola (STMicroelectronics)*
- 5446 *Blasi, Danilo (STMicroelectronics)*
- 5447 *Bois, Simone (STMicroelectronics)*
- 5448 *Brunschweiler, Andreas (CURRENT Technologies International)*
- 5449 *Casone, Luca (STMicroelectronics)*
- 5450 *Cassin-Delauriere, Agnes (Texas Instruments)*
- 5451 *Du, Shu (Texas Instruments)*
- 5452 *Escrhuela, Francisco (Ormazabal CURRENT)*
- 5453 *Estopiñan, Pedro (Atmel)*
- 5454 *Garai, Mikel (ZIV)*
- 5455 *Grasso, Riccardo (STMicroelectronics)*
- 5456 *Kehn, Doug (Ormazabal CURRENT)*
- 5457 *Kim, Il Han (Texas Instruments)*
- 5458 *Lasciandare, Alessandro (STMicroelectronics)*
- 5459 *Liu, Weilin (CURRENT Technologies International)*
- 5460 *Llano, Asier (ZIV)*
- 5461 *Llorente, Isabel (Naturgy)*
- 5462 *Lunn, Andrew (CURRENT Technologies International)*
- 5463 *Manero, Eduardo (Atmel)*
- 5464 *Muñoz, Andrés (Atmel)*
- 5465 *Pulkkinen, Anssi (CURRENT Technologies International)*
- 5466 *Rodriguez Roncero, Javier (Landis+Gyr)*

-
- 5467 *Romero, Gloria (Itron)*
 - 5468 *Saccani, Emile (ST Microelectronics)*
 - 5469 *Sánchez, Agustín (Landis+Gyr)*
 - 5470 *Sanz, Alfredo (Atmel)*
 - 5471 *Scarpa, Vincenzo (STMicroelectronics)*
 - 5472 *Schaub, Thomas (Landis+Gyr)*
 - 5473 *Sedjai, Mohamed (CURRENT Technologies International)*
 - 5474 *Sendin, Alberto (Iberdrola)*
 - 5475 *Sharma, Manu (CURRENT Technologies International)*
 - 5476 *Susella, Ruggero (STMicroelectronics)*
 - 5477 *Tarruell, Frederic (Itron)*
 - 5478 *Teijeiro, Jesús (Atmel)*
 - 5479 *Treffiletti, Paolo (STMicroelectronics)*
 - 5480 *Varadarajan, Badri (Texas Instruments)*
 - 5481 *Widmer, Hanspeter (CURRENT Technologies International)*
 - 5482 *Wikiera, Jacek (CURRENT Technologies International)*