

IEEE P802.1AX-REV™/D4.54

Draft Standard for
Local and Metropolitan Area Networks —

Link Aggregation

Sponsor

LAN MAN Standards Committee
of the
IEEE Computer Society

Prepared by the Interworking Task Group of IEEE 802.1

Abstract: Link Aggregation allows parallel full duplex point-to-point links to be used as if they were a single link, and also supports the use of multiple links as a resilient load sharing interconnect between multiple nodes in two separately administered networks. This standard defines a MAC independent Link Aggregation capability, and general information relevant to specific to MAC types.

Keywords: Aggregated Link, Aggregator, Link Aggregation, Link Aggregation Group, local area network, management, interconnect, Network-Network Interface, NNI, Distributed Resilient Network Interconnect, DRNI.

Copyright © 2014 by the Institute of Electrical and Electronics Engineers, Inc.
345 East 47th Street
New York, NY 10017, USA
All rights reserved.

All rights reserved. This document is an unapproved draft of a proposed IEEE Standard. As such, this document is subject to change. USE AT YOUR OWN RISK! Because this is an unapproved draft, this document must not be utilized for any conformance/compliance purposes. Permission is hereby granted for IEEE Standards Committee participants to reproduce this document for purposes of IEEE standardization activities only. Prior to submitting this document to another standards development organization for standardization activities, permission must first be obtained from the Manager, Standards Licensing and Contracts, IEEE Standards Activities Department. Other entities seeking permission to reproduce this document, in whole or in part, must obtain permission from the Manager, Standards Licensing and Contracts, IEEE Standards Activities Department.

IEEE Standards Department
Copyright and Permissions
445 Hoes Lane, P.O. Box 1331
Piscataway, NJ 08855-1331, USA

Introduction to IEEE P802.1AX-REV™

This introduction is not part of IEEE Project 802.1AX-REV, IEEE Standards for Local and Metropolitan Area Networks—Link Aggregation

This standard contains state-of-the-art material. The area covered by this standard is undergoing evolution. Revisions are anticipated within the next few years to clarify existing material, to correct possible errors, and to incorporate new related material. Information on the current revision state of this and other IEEE 802 standards can be obtained from

Secretary, IEEE-SA Standards Board
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
USA

Notice to Users

Laws and regulations

Users of these documents should consult all applicable laws and regulations. Compliance with the provisions of this standard does not imply compliance to any applicable regulatory requirements. Implementers of the standard are responsible for observing or referring to the applicable regulatory requirements. IEEE does not, by the publication of its standards, intend to urge action that is not in compliance with applicable laws, and these documents may not be construed as doing so.

Copyrights

This document is copyrighted by the IEEE. It is made available for a wide variety of both public and private uses. These include both use, by reference, in laws and regulations, and use in private self-regulation, standardization, and the promotion of engineering practices and methods. By making this document available for use and adoption by public authorities and private users, the IEEE does not waive any rights in copyright to this document.

Updating of IEEE documents

Users of IEEE standards should be aware that these documents may be superseded at any time by the issuance of new editions or may be amended from time to time through the issuance of amendments, corrigenda, or errata. An official IEEE document at any point in time consists of the current edition of the document together with any amendments, corrigenda, or errata then in effect. In order to determine whether a given document is the current edition and whether it has been amended through the issuance of amendments, corrigenda, or errata, visit the IEEE Standards Association website at <http://ieeexplore.ieee.org/xpl/standards.jsp>, or contact the IEEE at the address listed previously.

For more information about the IEEE Standards Association or the IEEE standards development process, visit the IEEE-SA website at <http://standards.ieee.org>.

Errata

Errata, if any, for this and all other standards can be accessed at the following URL: <http://standards.ieee.org/reading/ieee/updates/errata/index.html>. Users are encouraged to check this URL for errata periodically.

Interpretations

Current interpretations can be accessed at the following URL: <http://standards.ieee.org/reading/ieee/interp/index.html>.

Patents

Attention is called to the possibility that implementation of this amendment may require use of subject matter covered by patent rights. By publication of this amendment, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE is not responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patents Claims or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or nondiscriminatory.

Users of this amendment are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility. Further information may be obtained from the IEEE Standards Association.

Contents

1		
2		
3	1.	Overview 13
4	1.1	Scope 13
5	1.2	Purpose 13
6	1.3	State diagram conventions 13
7		
8	2.	Normative references 14
9		
10	3.	Definitions 15
11		
12	4.	Abbreviations 18
13		
14	5.1	Requirements Terminology 19
15	5.2	Protocol Implementation Conformance Statement (PICS) 19
16	5.3	Link Aggregation requirements 19
17	5.3.1	Link Aggregation options 19
18	5.4	Distributed Resilient Network Interconnect requirements 20
19	5.4.1	Distribution Resilient Network Interconnect options 20
20		
21	6.	Link Aggregation 22
22	6.1	Overview 22
23	6.1.1	Goals and objectives 22
24	6.1.2	Positioning of Link Aggregation within the IEEE 802 architecture 23
25	6.1.3	LLDP Parser/Multiplexer 24
26	6.1.3.1	LLDP Parser state diagram 24
27	6.1.3.1.1	LLDP Parser Function 24
28	6.1.3.1.2	Constants 24
29	6.1.3.1.3	Variables 25
30	6.1.3.1.4	State diagram 25
31	6.2	Link Aggregation operation 25
32	6.2.1	Principles of Link Aggregation 25
33	6.2.2	Service interfaces 27
34	6.2.3	Frame Collector 28
35	6.2.3.1	Frame Collector state diagram 28
36	6.2.3.1.1	Constants 28
37	6.2.3.1.2	Variables 28
38	6.2.3.1.3	Messages 28
39	6.2.3.1.4	State diagram 29
40	6.2.4	Frame Distributor 29
41	6.2.4.1	Frame Distributor state diagram 30
42	6.2.4.1.1	Variables 30
43	6.2.4.1.2	Messages 30
44	6.2.4.1.3	State diagram 31
45	6.2.5	Marker Generator/Receiver (optional) 31
46	6.2.6	Marker Responder 31
47	6.2.7	Protocol Parser/Multiplexer 32
48	6.2.7.1	Protocol Parser state diagram 32
49	6.2.7.1.1	Functions 32
50	6.2.7.1.2	Variables 32
51	6.2.7.1.3	Messages 32
52	6.2.7.1.4	State diagram 33
53	6.2.8	Aggregator Parser/Multiplexer 33
54	6.2.8.1	Aggregator Parser state diagram 33
	6.2.8.1.1	Constants 33
	6.2.8.1.2	Variables 34

1		6.2.8.1.3	Messages	34
2		6.2.8.1.4	State Diagram	35
3	6.2.9	Aggregator		35
4	6.2.10	Control Parser/Multiplexer		36
5		6.2.10.1	Control Parser state diagram	36
6		6.2.10.1.1	Control Parser Function	36
7		6.2.10.1.2	Constants	36
8		6.2.10.1.3	Variables	36
9	6.2.11	Addressing		37
10		6.2.11.1	Source Address	37
11		6.2.11.2	Destination Address	37
12	6.3	Link Aggregation Control		38
13		6.3.1	Characteristics of Link Aggregation Control	39
14		6.3.2	System identification	40
15		6.3.3	Aggregator identification	40
16		6.3.4	Port identification	40
17		6.3.5	Capability identification	41
18		6.3.6	Link Aggregation Group identification	42
19		6.3.6.1	Construction of the Link Aggregation Group Identifier	42
20		6.3.6.2	Representation of the Link Aggregation Group Identifier	43
21		6.3.7	Selecting a Link Aggregation Group	43
22		6.3.8	Agreeing on a Link Aggregation Group	44
23		6.3.9	Attaching a link to an Aggregator	44
24		6.3.10	Signaling readiness to transfer user data	44
25		6.3.11	Enabling the Frame Collector and Frame Distributor	45
26		6.3.12	MAC_Operational status	45
27		6.3.13	Monitoring the membership of a Link Aggregation Group	45
28		6.3.14	Detaching a link from an Aggregator	46
29		6.3.15	Configuration and administrative control of Link Aggregation	46
30		6.3.16	Link Aggregation Control state information	46
31	6.4	Link Aggregation Control Protocol (LACP)		47
32		6.4.1	LACP design elements	47
33		6.4.2	LACPDU structure and encoding	47
34		6.4.2.1	Transmission and representation of octets	47
35		6.4.2.2	Encapsulation of LACPDU in frames	48
36		6.4.2.3	LACPDU structure	48
37		6.4.2.4	Version 2 TLVs	52
38		6.4.2.4.1	Port Algorithm TLV	52
39		6.4.2.4.2	Port Conversation ID Digest TLV	53
40		6.4.2.4.3	Port Conversation Mask TLVs	53
41		6.4.2.4.4	Port Conversation Service Mapping TLV	56
42		6.4.3	LACP state machine overview	56
43		6.4.4	Constants	57
44		6.4.5	Variables associated with the System	59
45		6.4.6	Variables associated with each Aggregator	59
46		6.4.7	Variables associated with each Aggregation Port	60
47		6.4.8	Variables used for managing the operation of the state machines	63
48		6.4.9	Functions	64
49		6.4.10	Timers	66
50		6.4.11	Messages	66
51		6.4.12	Receive machine	66
52		6.4.13	Periodic Transmission machine	68
53		6.4.14	Selection Logic	70
54		6.4.14.1	Selection Logic—Requirements	70
		6.4.14.2	Selection Logic—Recommended default operation	72
		6.4.15	Mux machine	73
		6.4.16	Transmit machine	76
		6.4.17	Churn Detection machines	77

1	6.4.18	Long LACPDU machine	77
2	6.5	Marker protocol	79
3	6.5.1	Introduction	79
4	6.5.2	Sequence of operations	80
5	6.5.3	Marker and Marker Response PDU structure and encoding	80
6	6.5.3.1	Transmission and representation of octets	80
7	6.5.3.2	Encapsulation of Marker and Marker Response PDU in frames	80
8	6.5.3.3	Marker and Marker Response PDU structure	81
9	6.5.4	Protocol definition	82
10	6.5.4.1	Operation of the marker protocol	82
11	6.5.4.2	Marker Responder state diagram	83
12	6.5.4.2.1	Variables	83
13	6.5.4.2.2	Messages	83
14	6.6	Conversation-sensitive frame collection and distribution	84
15	6.6.1	Conversation-sensitive collection and distribution state diagrams	84
16	6.6.1.1	Conversion-sensitive collection state diagram	85
17	6.6.1.1.1	Variables	85
18	6.6.1.1.2	Variables associated with each Aggregation Port	85
19	6.6.1.1.3	Functions	85
20	6.6.1.1.4	Messages	85
21	6.6.1.1.5	State diagram	86
22	6.6.2	Conversation-sensitive LACP state diagrams	86
23	6.6.2.1	Per-Aggregator Variables	87
24	6.6.2.2	Variables associated with each Aggregation Port	88
25	6.6.2.3	Variables used for managing the operation of the state diagrams	90
26	6.6.2.4	Functions	90
27	6.6.2.5	Timers	93
28	6.6.2.6	Messages	93
29	6.6.2.7	State diagrams	93
30	6.7	Configuration capabilities and restrictions	99
31	6.7.1	Use of system and port priorities	99
32	6.7.2	Dynamic allocation of operational Keys	99
33	6.7.3	Link Aggregation on shared-medium links	100
34	6.7.4	Selection Logic variants	100
35	6.7.4.1	Reduced reconfiguration	101
36	6.7.4.2	Limited Aggregator availability	101
37	7.	Management	102
38	7.1	Overview	102
39	7.1.1	Systems management overview	102
40	7.1.2	Management model	103
41	7.2	Managed objects	103
42	7.2.1	Introduction	103
43	7.2.2	Overview of managed objects	104
44	7.2.2.1	Text description of managed objects	104
45	7.2.3	Containment	105
46	7.2.4	Naming	105
47	7.2.5	Capabilities	105
48	7.3	Management for Link Aggregation	110
49	7.3.1	Aggregator managed object class	110
50	7.3.1.1	Aggregator attributes	111
51	7.3.1.1.1	aAggID	111
52	7.3.1.1.2	aAggDescription	111
53	7.3.1.1.3	aAggName	112
54	7.3.1.1.4	aAggActorSystemID	112
	7.3.1.1.5	aAggActorSystemPriority	112
	7.3.1.1.6	aAggAggregateOrIndividual	112

1	7.3.1.1.7	aAggActorAdminKey	112
2	7.3.1.1.8	aAggActorOperKey	113
3	7.3.1.1.9	aAggMACAddress	113
4	7.3.1.1.10	aAggPartnerSystemID	113
5	7.3.1.1.11	aAggPartnerSystemPriority	113
6	7.3.1.1.12	aAggPartnerOperKey	113
7	7.3.1.1.13	aAggAdminState	114
8	7.3.1.1.14	aAggOperState	114
9	7.3.1.1.15	aAggTimeOfLastOperChange	114
10	7.3.1.1.16	aAggDataRate	114
11	7.3.1.1.17	aAggOctetsTxOK	115
12	7.3.1.1.18	aAggOctetsRxOK	115
13	7.3.1.1.19	aAggFramesTxOK	115
14	7.3.1.1.20	aAggFramesRxOK	115
15	7.3.1.1.21	aAggMulticastFramesTxOK	116
16	7.3.1.1.22	aAggMulticastFramesRxOK	116
17	7.3.1.1.23	aAggBroadcastFramesTxOK	116
18	7.3.1.1.24	aAggBroadcastFramesRxOK	116
19	7.3.1.1.25	aAggFramesDiscardedOnTx	117
20	7.3.1.1.26	aAggFramesDiscardedOnRx	117
21	7.3.1.1.27	aAggFramesWithTxErrors	117
22	7.3.1.1.28	aAggFramesWithRxErrors	117
23	7.3.1.1.29	aAggUnknownProtocolFrames	118
24	7.3.1.1.30	aAggPortList	118
25	7.3.1.1.31	aAggLinkUpDownNotificationEnable	118
26	7.3.1.1.32	aAggCollectorMaxDelay	118
27	7.3.1.1.33	aAggPortAlgorithm	118
28	7.3.1.1.34	aAggPartnerAdminPortAlgorithm	119
29	7.3.1.1.35	aAggConversationAdminLink[]	119
30	7.3.1.1.36	aAggPartnerAdminPortConversationListDigest	119
31	7.3.1.1.37	aAggAdminDiscardWrongConversation	120
32	7.3.1.1.38	aAggAdminServiceConversationMap[]	120
33	7.3.1.1.39	aAggPartnerAdminConvServiceMappingDigest	120
34	7.3.1.2	Aggregator Notifications	120
35	7.3.1.2.1	nAggLinkUpNotification	120
36	7.3.1.2.2	nAggLinkDownNotification	121
37	7.3.2	Aggregation Port managed object class	121
38	7.3.2.1	Aggregation Port Attributes	121
39	7.3.2.1.1	aAggPortID	121
40	7.3.2.1.2	aAggPortActorSystemPriority	121
41	7.3.2.1.3	aAggPortActorSystemID	121
42	7.3.2.1.4	aAggPortActorAdminKey	121
43	7.3.2.1.5	aAggPortActorOperKey	122
44	7.3.2.1.6	aAggPortPartnerAdminSystemPriority	122
45	7.3.2.1.7	aAggPortPartnerOperSystemPriority	122
46	7.3.2.1.8	aAggPortPartnerAdminSystemID	122
47	7.3.2.1.9	aAggPortPartnerOperSystemID	122
48	7.3.2.1.10	aAggPortPartnerAdminKey	123
49	7.3.2.1.11	aAggPortPartnerOperKey	123
50	7.3.2.1.12	aAggPortSelectedAggID	123
51	7.3.2.1.13	aAggPortAttachedAggID	123
52	7.3.2.1.14	aAggPortActorPort	124
53	7.3.2.1.15	aAggPortActorPortPriority	124
54	7.3.2.1.16	aAggPortPartnerAdminPort	124
	7.3.2.1.17	aAggPortPartnerOperPort	124
	7.3.2.1.18	aAggPortPartnerAdminPortPriority	124
	7.3.2.1.19	aAggPortPartnerOperPortPriority	125
	7.3.2.1.20	aAggPortActorAdminState	125
	7.3.2.1.21	aAggPortActorOperState	125

1		7.3.2.1.22	aAggPortPartnerAdminState	125
2		7.3.2.1.23	aAggPortPartnerOperState	126
3		7.3.2.1.24	aAggPortAggregateOrIndividual	126
4		7.3.2.1.25	aAggPortOperConversationPasses	126
5		7.3.2.1.26	aAggPortOperConversationCollected	126
6		7.3.2.1.27	aAggPortLinkNumberID	126
7		7.3.2.1.28	aAggPortPartnerAdminLinkNumberID	127
8		7.3.2.1.29	aAggPortWTRTime	127
9		7.3.2.2	Aggregation Port Extension Attributes	127
10		7.3.2.2.1	aAggPortProtocolDA	127
11	7.3.3	Aggregation Port Statistics managed object class	127	
12		7.3.3.1	Aggregation Port Statistics attributes	128
13		7.3.3.1.1	aAggPortStatsID	128
14		7.3.3.1.2	aAggPortStatsLACPDUsRx	128
15		7.3.3.1.3	aAggPortStatsMarkerPDUsRx	128
16		7.3.3.1.4	aAggPortStatsMarkerResponsePDUsRx	128
17		7.3.3.1.5	aAggPortStatsUnknownRx	128
18		7.3.3.1.6	aAggPortStatsIllegalRx	129
19		7.3.3.1.7	aAggPortStatsLACPDUsTx	129
20		7.3.3.1.8	aAggPortStatsMarkerPDUsTx	129
21		7.3.3.1.9	aAggPortStatsMarkerResponsePDUsTx	129
22	7.3.4	Aggregation Port Debug Information managed object class	129	
23		7.3.4.1	Aggregation Port Debug Information attributes	129
24		7.3.4.1.1	aAggPortDebugInformationID	129
25		7.3.4.1.2	aAggPortDebugRxState	130
26		7.3.4.1.3	aAggPortDebugLastRxTime	130
27		7.3.4.1.4	aAggPortDebugMuxState	130
28		7.3.4.1.5	aAggPortDebugMuxReason	131
29		7.3.4.1.6	aAggPortDebugActorChurnState	131
30		7.3.4.1.7	aAggPortDebugPartnerChurnState	131
31		7.3.4.1.8	aAggPortDebugActorChurnCount	131
32		7.3.4.1.9	aAggPortDebugPartnerChurnCount	132
33		7.3.4.1.10	aAggPortDebugActorSyncTransitionCount	132
34		7.3.4.1.11	aAggPortDebugPartnerSyncTransitionCount	132
35		7.3.4.1.12	aAggPortDebugActorChangeCount	132
36		7.3.4.1.13	aAggPortDebugPartnerChangeCount	132
37		7.3.4.1.14	aAggPortDebugActorCDSChurnState	132
38		7.3.4.1.15	aAggPortDebugPartnerCDSChurnState	133
39		7.3.4.1.16	aAggPortDebugActorCDSChurnCount	133
40		7.3.4.1.17	aAggPortDebugPartnerCDSChurnCount	133
41	7.4	Management for Distributed Resilient Network Interconnect	133	
42		7.4.1	Distributed Relay Managed Object Class	133
43		7.4.1.1	Distributed Relay Attributes	134
44		7.4.1.1.1	aDrmiID	134
45		7.4.1.1.2	aDrmiDescription	134
46		7.4.1.1.3	aDrmiName	134
47		7.4.1.1.4	aDrmiPortalAddr	134
48		7.4.1.1.5	aDrmiPortalPriority	134
49		7.4.1.1.6	aDrmiThreePortalSystem	135
50		7.4.1.1.7	aDrmiPortalSystemNumber	135
51		7.4.1.1.8	aDrmiIntraPortalLinkList	135
52		7.4.1.1.9	aDrmiAggregator	135
53		7.4.1.1.10	aDrmiConvAdminGateway[]	135
54		7.4.1.1.11	aDrmiNeighborAdminConvGatewayListDigest	136
		7.4.1.1.12	aDrmiNeighborAdminConvPortListDigest	136
		7.4.1.1.13	aDrmiGatewayAlgorithm	136
		7.4.1.1.14	aDrmiNeighborAdminGatewayAlgorithm	136
		7.4.1.1.15	aDrmiNeighborAdminPortAlgorithm	137
		7.4.1.1.16	aDrmiNeighborAdminDRCPState	137

1		7.4.1.1.17	aDrniEncapsulationMethod	137
2		7.4.1.1.18	aDrniIPLEncapMap	137
3		7.4.1.1.19	aDrniNetEncapMap	138
4		7.4.1.1.20	aDrniDRPortConversationPasses	138
5		7.4.1.1.21	aDrniDRGatewayConversationPasses	138
6		7.4.1.1.22	aDrniPSI	138
7		7.4.1.1.23	aDrniPortConversationControl	139
8		7.4.1.1.24	aDrniIntraPortalPortProtocolDA	139
9	7.4.2		IPP Managed Objects Class	139
10		7.4.2.1	IPP Attributes	139
11		7.4.2.1.1	aIPPID	139
12		7.4.2.1.2	aIPPortConversationPasses	139
13		7.4.2.1.3	aIPPGatewayConversationDirection	140
14		7.4.2.1.4	aIPAdminState	140
15		7.4.2.1.5	aIPOperState	140
16		7.4.2.1.6	aIPTimeOfLastOperChange	140
17	7.4.3		IPP Statistics managed object class	141
18		7.4.3.1	IPP Statistics attributes	141
19		7.4.3.1.1	aIPStatsID	141
20		7.4.3.1.2	aIPStatsDRCPDUsRx	141
21		7.4.3.1.3	aIPStatsIllegalRx	141
22		7.4.3.1.4	aIPStatsDRCPDUsTx	141
23	7.4.4		IPP Debug Information managed object class	141
24		7.4.4.1	IPP Debug Information attributes	142
25		7.4.4.1.1	aIPDebugInformationID	142
26		7.4.4.1.2	aIPDebugDRCPRxState	142
27		7.4.4.1.3	aIPDebugLastRxTime	142
28		7.4.4.1.4	aIPDebugDifferPortalReason	142
29	8.		Frame distribution and collection algorithms	143
30		8.1	Conversation Identifiers	143
31		8.2	Per-service frame distribution	143
32		8.2.1	Goals and objectives	143
33		8.2.2	Overview	143
34		8.2.3	Port Conversation Identifiers	144
35	9.		Distributed Resilient Network Interconnect	145
36		9.1	Goals and Objectives	145
37		9.2	Distributed Relay	146
38		9.3	Distributed Relay operation and procedures	148
39		9.3.1	Portal Topology	151
40		9.3.2	Intra-Portal Link	152
41		9.3.2.1	Network / IPL sharing by time	152
42		9.3.2.2	Network / IPL sharing by tag	153
43		9.3.2.3	Network / IPL sharing by encapsulation	153
44		9.3.3	Protocol Identification	154
45		9.3.4	DR Function state machines	154
46		9.3.4.1	Service interfaces	155
47		9.3.4.2	Per-DR Function variables	155
48		9.3.4.3	Per-IPP variables	156
49		9.3.4.4	Functions	156
50		9.3.4.5	Messages	157
51		9.3.4.6	DR Function: Aggregator Port reception state machine	157
52		9.3.4.7	DR Function: Gateway distribution state machine	157
53		9.3.4.8	DR Function: IPP N reception state machine	158
54	9.4		Distributed Relay Control Protocol	159
		9.4.1	Establishing the Portal and Distributed Relay	161

1	9.4.2	DRCPDU transmission, addressing, and protocol identification	161
2	9.4.2.1	Destination MAC Address	162
3	9.4.2.2	Source MAC Address	162
4	9.4.2.3	Priority	162
5	9.4.2.4	Encapsulation of DRCPDUs in frames	162
6	9.4.3	DRCPDU structure and encoding	162
7	9.4.3.1	Transmission and representation of octets	162
8	9.4.3.2	DRCPDU structure	163
9	9.4.3.3	Conversation Vector TLVs	170
10	9.4.3.3.1	2P Gateway Conversation Vector TLV	170
11	9.4.3.3.2	3P Gateway Conversation Vector-1 TLV	171
12	9.4.3.3.3	3P Gateway Conversation Vector-2 TLV	171
13	9.4.3.3.4	2P Port Conversation Vector TLV	172
14	9.4.3.3.5	3P Port Conversation Vector-1 TLV	173
15	9.4.3.3.6	3P Port Conversation Vector-2 TLV	173
16	9.4.3.4	Network/IPL sharing TLVs	174
17	9.4.3.4.1	Network/IPL Sharing Method TLV	174
18	9.4.3.4.2	Network/IPL Sharing Encapsulation TLV	175
19	9.4.3.5	Organization-Specific TLV	175
20	9.4.4	DRCP Control Parser/Multiplexer	176
21	9.4.4.1	Control Parser state diagram	176
22	9.4.4.1.1	Control Parser Function	176
23	9.4.4.1.2	Constants	177
24	9.4.4.1.3	Variables	177
25	9.4.5	DRCP state machine overview	177
26	9.4.6	Constants	178
27	9.4.7	Variables associated with the Distributed Relay	178
28	9.4.8	Per-DR Function variables	180
29	9.4.9	Per-IPP variables	183
30	9.4.10	Variables used for managing the operation of the state machines	189
31	9.4.11	Functions	190
32	9.4.12	Timers	204
33	9.4.13	Messages	204
34	9.4.14	DRCPDU Receive machine	204
35	9.4.15	DRCP Periodic Transmission machine	206
36	9.4.16	Portal System machine	208
37	9.4.17	DRNI Gateway and Aggregator machines	209
38	9.4.18	DRNI IPP machines	210
39	9.4.19	DRCPDU Transmit machine	211
40	9.4.20	Network/IPL sharing machine	212
41	Annex A (normative)	Protocol Implementation Conformance Statement (PICS) proforma	214
42	A.1	Introduction	214
43	A.1.1	Abbreviations and special symbols	214
44	A.1.2	Instructions for completing the PICS proforma	215
45	A.1.3	Additional information	215
46	A.1.4	Exceptional information	215
47	A.1.5	Conditional items	216
48	A.1.6	Identification	216
49	A.2	PICS proforma for Clause 6, Link Aggregation	217
50	A.2.1	Major capabilities/options	217
51	A.2.2	LLDP Port connectivity	218
52	A.2.3	Protocol Parser/Multiplexer support	218
53	A.2.4	Frame Collector	218
54	A.2.5	Frame Distributor	218
55	A.2.6	Marker protocol	219
56	A.2.7	Aggregator Parser/Multiplexer	219
57	A.2.8	Control Parser/Multiplexer	220

1	A.2.9	System identification	220
2	A.2.10	Aggregator identification	220
3	A.2.11	Port identification	221
4	A.2.12	Capability identification	221
5	A.2.13	Link Aggregation Group identification	221
6	A.2.14	Detaching a link from an Aggregator	221
7	A.2.15	LACPDU structure	222
8	A.2.16	Version 2 LACPDU	222
9	A.2.17	State machine variables	222
10	A.2.18	Receive machine	223
11	A.2.19	Periodic Transmission machine	223
12	A.2.20	Selection Logic	223
13	A.2.21	Mux machine	224
14	A.2.22	Transmit machine	224
15	A.2.23	Churn Detection machines	225
16	A.2.24	Marker protocol	225
17	A.2.25	Management	227
18	A.2.26	Per-Service Frame Distribution	230
19	A.2.27	Conversation-sensitive frame collection and distribution	230
20	A.2.28	Configuration capabilities and restrictions	231
21	A.2.29	Link Aggregation on shared-medium links	231
22	A.2.30	Distributed Resilient Network Interconnect	232
23	A.2.31	DRCPDU structure	233
24	A.2.32	Bridge specific support	234
25	Annex B (normative)	Collection and distribution algorithms	235
26	B.1	Introduction	235
27	B.2	Port selection	236
28	B.3	Dynamic reallocation of conversations to different Aggregation Ports	236
29	B.4	Topology considerations in the choice of distribution algorithm	237
30	Annex C (normative)	LACP standby link selection and dynamic Key management.....	239
31	C.1	Introduction	239
32	C.2	Goals	239
33	C.3	Standby link selection	240
34	C.4	Dynamic Key management	240
35	C.5	A dynamic Key management algorithm	240
36	C.6	Example 1	241
37	C.7	Example 2	242
38	Annex D (normative)	SMIv2 MIB definitions for Link Aggregation.....	244
39	D.1	Introduction	244
40	D.2	The SNMP Management Framework	244
41	D.3	Security considerations	244
42	D.4	Structure of the MIB module	245
43	D.4.1	Relationship to the managed objects defined in Clause 7	246
44	D.4.2	MIB Subtrees	249
45	D.5	Relationship to other MIBs	251
46	D.5.1	Relationship to the Interfaces MIB	251
47	D.5.2	Layering model	252
48	D.5.3	ifStackTable	252
49	D.5.4	ifRcvAddressTable	252
50	D.6	Definitions for Link Aggregation MIB	252
51	Annex E (normative)	Distributed Bridge.....	319
52			
53			
54			

1	E.1	Distributed VLAN Bridge	319
2	E.2	Higher Layer Entities in a Distributed Bridge	323
3			
4	Annex F (normative)	Link Layer Discovery Protocol TLVs	324
5	F.1	Link Aggregation TLV	324
6	F.1.1	aggregation status	324
7	F.1.2	aggregated Port ID	325
8	F.1.3	Link Aggregation TLV usage rules	325
9	F.1.4	Use of other TLVs on an Aggregator or Aggregation Link	325
10			
11	Annex G (normative)	Network / IPL sharing by time - MAC Address synchronization	327
12	G.1	Address Synchronization - design goals	328
13	G.2	Address Synchronization - non-goals.	328
14	G.3	Protocol Summary	328
15	G.4	Address Synchronization Description	328
16	G.5	ASPDU transmission, addressing, and protocol identification	330
17	G.5.1	Destination MAC Address	330
18	G.5.2	Source MAC Address	330
19	G.5.3	Priority	330
20	G.5.4	Encapsulation of ASPDUs in frames	330
21	G.5.5	ASPDU structure and encoding	331
22			
23			
24			
25			
26			
27			
28			
29			
30			
31			
32			
33			
34			
35			
36			
37			
38			
39			
40			
41			
42			
43			
44			
45			
46			
47			
48			
49			
50			
51			
52			
53			
54			

IEEE P802.1AX-REV™/D4.54

Draft Standard for Local and Metropolitan Area Networks— Link Aggregation

1. Overview

1.1 Scope

Link Aggregation provides protocols, procedures, and managed objects that allow:

- One or more parallel instances of full duplex point-to-point links to be aggregated together to form a Link Aggregation Group, such that a MAC Client can treat the Link Aggregation Group as if it were a single link.
- A resilient interconnect using multiple full duplex point-to-point links among one to three nodes in a network and one to three nodes in another, separately administered, network, along with a means to ensure that frames belonging to any given service will use the same physical path in both directions between the two networks.

This standard defines the MAC independent Link Aggregation capability, and general information relevant to specific MAC types that support Link Aggregation. The capabilities defined are compatible with previous versions of this standard.

1.2 Purpose

Link Aggregation allows the establishment of full duplex point-to-point links that have a higher aggregate bandwidth than the individual links that form the aggregation, and the use of multiple systems at each end of the aggregation. This allows improved utilization of available links in bridged LAN environments, along with improved resilience in the face of failure of individual links or systems. In applications connecting separately administered networks, the networks are isolated from each other's fault recovery events.

1.3 State diagram conventions

This document uses the state diagram conventions of IEEE Std 802.1Q-2011, Annex E.

Should a conflict exist between a state diagram and either the corresponding global transition tables or the textual description associated with the state machine, the state diagram takes precedence.

2. Normative references

The following referenced documents are indispensable for the application of this document (i.e., they are to be understood and used, so each referenced document is cited in text and its relationship to this document is explained). For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments or corrigenda) applies.

IEEE Std 802[®], IEEE Standards for Local and Metropolitan Area Networks: Overview and Architecture.^{1, 2}

IEEE Std 802.1AC[™], IEEE Standard for Local and Metropolitan Area Networks Media Access Control (MAC) Service Definition.

IEEE Std 802.1D[™], IEEE Standard for Local and metropolitan area networks Media Access Control (MAC) Bridges.

IEEE Std 802.1Q[™]-2011, IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks.

IEEE Std 802.3[™]-2012, IEEE Standard for Ethernet.

IETF RFC 1213 (IETF STD 17), *Management Information Base for Network Management of TCP/IP-based internets: MIB-II*, McCloghrie K., and M. Rose, Editors, March 1991³.

IETF RFC 1321, *The MD5 Message-Digest Algorithm*, R. Rivest. April 1992.

IETF RFC 2578 (STD 58) *Structure of Management Information Version 2 (SMIv2)*, K. McCloghrie, D. Perkins, J. Schoenwaelder. April 1999.

IETF RFC 2579 (STD 58) *Textual Conventions for SMIv2*, McCloghrie, K., Perkins, D., Schoenwaelder, J., Case, J., Rose, M. and S. Waldbusser, April 1999.

IETF RFC 2580 (STD 58) *Conformance Statements for SMIv2*, McCloghrie, K., Perkins, D., Schoenwaelder, J., Case, J., Rose, M. and S. Waldbusser, April 1999.

IETF RFC 2863, *The Interfaces Group MIB*, K. McCloghrie, F. Kastenholz, June 2000.

IETF RFC 3410, *Introduction and Applicability Statements for Internet-Standard Management Framework*, J. Case, R. Mundy, D. Partain, B. Stewart. December 2002.

IETF RFC 3414 (STD 62) *User-based Security Model (USM) for Version 3 of the Simple Network Management Protocol (SNMPv3)*, U. Blumenthal, B. Wijnen. December 2002.

IETF RFC 3415 (STD 62) *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*, B. Wijnen, R. Presuhn, K. McCloghrie, December 2002.

ISO/IEC 10165-4: 1992, Information technology—Open Systems Interconnection—Management information services—Structure of management information—Part 4: Guidelines for the definition of managed objects.

1. IEEE publications are available from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, Piscataway, NJ, 08854, USA (<http://standards.ieee.org/>).

2. The IEEE standards or products referred to in this clause are trademarks of the Institute of Electrical and Electronics Engineers, Inc.

3. IETF RFCs are available from the Internet Engineering Task Force website at <http://www.ietf.org/rfc.html>.

3. Definitions

For the purposes of this standard, the following terms and definitions apply. *The Authoritative Dictionary of IEEE Standard Terms* (IEEE 100) should be referenced for terms not defined in this standard.

3.1 Actor: The local entity in a Link Aggregation Control Protocol exchange.

3.2 agent: A term used to refer to the managed nodes in a network. Managed nodes are those nodes that contain a network management entity, which can be used to configure the node and/or collect data describing operation of that node. The agent is controlled by a network control host or manager that contains both a network management entity and network management application software to control the operations of agents. Agents include Systems that support user applications as well as nodes that provide communications services such as front-end processors, bridges, and routers.

3.3 Aggregation Key: A parameter associated with each Aggregation Port and with each Aggregator of an Aggregation System identifying those Aggregation Ports that can be aggregated together. Aggregation Ports in an Aggregation System that share the same Aggregation Key value are potentially able to aggregate together.

3.4 Aggregation Link: A LAN connecting a pair of Aggregation Systems.

3.5 Aggregation Port: A Service Access Point (SAP) in an Aggregation System that is supported by a single MAC entity.

3.6 Aggregation System: A uniquely identifiable entity comprising (among other things) an arbitrary grouping of one or more Aggregation Ports for the purpose of aggregation. An instance of an aggregated link always occurs between exactly two Aggregation Systems. A physical device may comprise a single Aggregation System or more than one Aggregation System.

3.7 Aggregator: A logical MAC, bound to one or more Aggregation Ports, through which the Aggregator Client is provided access to the physical media.

3.8 Aggregator Client: The layered entity immediately above the Link Aggregation Sublayer, for which the Link Aggregation sublayer provides an instance of the Internal Sublayer Service (ISS).

3.9 Aggregator Port: A Service Access Point (SAP) in an Aggregation System that is supported by an Aggregator.

3.10 bridge: A layer 2 interconnection device that conforms to IEEE Std 802.1D and/or IEEE Std 802.1Q.

3.11 conversation: A set of frames transmitted from one end station to another, where all of the frames form an ordered sequence, and where the communicating end stations require the ordering to be maintained among the set of frames exchanged.

3.12 Conversation ID: An identifier using values in the range of 0 through 4095 to identify conversations.

3.13 data terminal equipment (DTE): Any source or destination of data connected to the local area network.

3.14 Distributed Relay: A functional entity, distributed over a Portal by a DR Function in each of the Systems comprising a Portal, that distributes down frames from Gateways to Aggregators, and distributes up frames from Aggregators to Gateways.

1 **3.15 Distributed Resilient Network Interconnect (DRNI):** Link Aggregation, as specified in Clause 6,
2 expanded to include at least one end of the Link Aggregation Group attached to a Portal such that
3 independence and isolation is provided between the interconnected network segments.
4

5 **3.16 Down frame:** A frame entering a Portal through a Gateway.
6

7 **3.17 DR Function:** That part of a Distributed Relay residing within a single Portal System.
8

9 **3.18 end station:** A System attached to a LAN that is an initial source or a final destination of frames
10 transmitted across that LAN.
11

12 **3.19 frame:** A unit of data transmission on a LAN that conveys a MAC Protocol Data Unit (MPDU) and
13 can cause a service indication with, at a minimum, a Destination Address, Source Address, a MAC Service
14 Data Unit (MSDU), and priority, or an MPDU that is the result of a service request with those parameters.
15

16 **3.20 full duplex:** A mode of operation of a network, DTE, or Medium Attachment Unit (MAU) that
17 supports duplex transmission as defined in IEEE 100. Within the scope of this standard, this mode of
18 operation allows for simultaneous communication between a pair of stations, provided that the Physical
19 Layer is capable of supporting simultaneous transmission and reception without interference. (See IEEE Std
20 802.3.)
21

22 **3.21 Gateway:** A connection, always virtual (not a physical link between Systems) connecting a
23 Distributed Relay to a System, consisting of a Gateway Link and two Gateway Ports.
24

25 **3.22 Gateway Conversation ID:** The Conversation ID value which is used within the Distributed Relay to
26 identify frames passing through a Gateway.
27

28 **3.23 Gateway Vector:** The operational Boolean vector, indexed by Gateway Conversation ID, indicating
29 whether the indexed Gateway Conversation ID is enabled to pass through a Portal System's Gateway
30 (FALSE = blocked). There is one Gateway Vector per Portal System in a Portal.
31

32 **3.24 half duplex:** A mode of operation of an CSMA/CD local area network (LAN) in which DTEs contend
33 for access to a shared medium. Multiple, simultaneous transmissions in a half duplex mode CSMA/CD LAN
34 result in interference, requiring resolution by the Ethernet protocol. (See IEEE Std 802.3.)
35

36 **3.25 Internal Sublayer Service (ISS):** An augmented version of the MAC Service, defined in IEEE Std
37 802.1AC.
38

39 **3.26 Intra-Portal Link (IPL):** A logical link used to connect the DR Functions comprising a Distributed
40 Relay.
41

42 **3.27 Intra-Portal Port (IPP):** A port to an Intra-Portal Link.
43

44 **3.28 Key:** *See:* Aggregation Key.
45

46 **3.29 link:** *See:* Aggregation Link.
47

48 **3.30 Link Aggregation Group:** A group of links that appear to an Aggregator Client as if they were a
49 single link. A Link Aggregation Group can connect two Aggregation Systems, an Aggregation System and a
50 Portal, or two Portals. One or more conversations may be associated with each link that is part of a Link
51 Aggregation Group.
52

53 **3.31 Long LACPDU:** A Version 2 or higher version LACPDU of a frame size that is larger than 128
54 octets.

1 **3.32 Management Information Base (MIB):** A repository of information to describe the operation of a
2 specific network device.
3

4 **3.33 Media Access Control (MAC):** The data link sublayer that is responsible for transferring data to and
5 from the Physical Layer.
6

7 **3.34 Partner:** The remote entity in a Link Aggregation Control Protocol exchange.
8

9 **3.35 Port Conversation ID:** The Conversation ID value that is used to select frames passing through an
10 Aggregation Port.
11

12 **3.36 Portal:** One end of a DRNI; One to three Systems, each with physical links that together comprise a
13 Link Aggregation Group. The Portal's Systems cooperate to emulate the presence of a single Aggregation
14 System to which the entire Link Aggregation Group is attached.
15

16 **3.37 Portal System:** A System that participates in a Portal.
17

18 **3.38 Portal System Number:** An integer from 1 through 3, inclusive, uniquely identifying a Portal System
19 within its Portal.
20

21 **3.39 selection algorithm:** The algorithm used to assign frames to Conversation IDs and Conversation IDs
22 to Aggregation Ports and, in the context of DRNI, also to Gateways.
23

24 **3.40 Service ID:** A value extracted from the header of a frame (VID, I-SID, etc.), received from an
25 Aggregation Link, that identifies the service instance with which that frame is associated.
26

27 **3.41 service instance:** A service instance is a set of Service Access Points (SAPs) such that a Data.Request
28 primitive presented to one SAP can result in a Data.Indication primitive occurring at one or more of the
29 other SAPs in that set. In the context of operators and customers, a particular customer is given access to all
30 of the SAPs of such a set by the operator.
31

32 **3.42 System:** *See:* Aggregation System.
33

34 **3.43 Type/Length/Value (TLV):** A variable length encoding of an information element consisting of
35 sequential type, length, and value fields where the type field identifies the type of information, the length
36 field indicates the length of the information field in octets, and the value field contains the information,
37 itself. The type value is locally defined and needs to be unique within the protocol defined in this standard.
38

39 **3.44 Up frame:** A frame entering a Portal through an Aggregation Port.
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

4. Abbreviations

This standard contains the following abbreviations:

ANSI	American National Standards Institute
CID	Company ID ¹ (See also OUI)
DA	destination address
DRCP	Distributed Relay Control Protocol
DRCPDU	Distributed Relay Control Protocol Data Unit
DRNI	Distributed Resilient Network Interconnect
DTE	data terminal equipment
FCS	frame check sequence
IEC	International Electrotechnical Commission
IPL	Intra-Portal Link
IPP	Intra-Portal Port
ISO	International Organization for Standardization
ISS	Internal Sublayer Service
LACP	Link Aggregation Control Protocol
LACPDU	Link Aggregation Control Protocol Data Unit
LAG	Link Aggregation Group
LAG ID	Link Aggregation Group Identifier
LAN	local area network
LLC	logical link control
LLID	logical link identifier
MAC	medium access control
MAN	metropolitan area network
MIB	management information base
NTT	Need To Transmit
OUI	organizationally unique identifier
PDU	Protocol Data Unit
PICS	protocol implementation conformance statement
PSI	Portal System Isolated
SA	source address
TLV	Type/Length/Value
TPMR	Two-Port Media Access Control (MAC) Relay
UCT	unconditional transition

1. See <http://standards.ieee.org/develop/regauth/tut/eui.pdf>

5. Conformance

This clause specifies the mandatory and optional capabilities provided by conformant implementations of this standard.

5.1 Requirements Terminology

For consistency with existing IEEE and IEEE 802.1 standards, requirements placed upon conformant implementations of this standard are expressed using the following terminology:

- a) **Shall** is used for mandatory requirements;
- b) **May** is used to describe implementation or administrative choices (“may” means “is permitted to,” and hence, “may” and “may not” mean precisely the same thing);
- c) **Should** is used for recommended choices (the behaviors described by “should” and “should not” are both permissible but not equally desirable choices).

The PICS proformas (see Annex A) reflect the occurrences of the words “shall,” “may,” and “should” within the standard.

The standard avoids needless repetition and apparent duplication of its formal requirements by using *is*, *is not*, *are*, and *are not* for definitions and the logical consequences of conformant behavior. Behavior that is permitted but is neither always required nor directly controlled by an implementer or administrator, or whose conformance requirement is detailed elsewhere, is described by *can*. Behavior that never occurs in a conformant implementation or System of conformant implementations is described by *cannot*. The word *allow* is used as a replacement for the phrase “Support the ability for,” and the word *capability* means “can be configured to.”

In this clause the term *conformant System* refers to a System that, for all ports for which support is claimed, conforms to the provisions of this standard for Link Aggregation (5.3) or Distributed Resilient Network Interconnect (5.4).

5.2 Protocol Implementation Conformance Statement (PICS)

The supplier of an implementation that is claimed to conform to this standard shall provide the information necessary to identify both the supplier and the implementation, and shall complete a copy of the PICS proforma provided in Annex A.

5.3 Link Aggregation requirements

A conformant System, shall:

- a) Support the Link Aggregation Sublayer and conform to the state machines and procedures in 6.2;
- b) Support the Link Aggregation Control Protocol and conform to the state machines and procedures in 6.3 and 6.4;
- c) Transmit and receive version 1 LACPDUs in the formats specified in 6.4.2;
- d) Implement the Marker Responder as specified in 6.5.4.2 and respond to received Marker PDUs as specified in 6.5.1.

5.3.1 Link Aggregation options

A conformant System, may:

- a) Support the Marker Protocol and conform to the state machines and procedures in 6.5 and transmit and receive required Marker and Marker Response PDUs in the formats specified in 6.5.3;
- b) Support the management functionality for Link Aggregation as specified in Clause 7;
- c) Support SMIPv2 MIB modules for the management of Link Aggregation capabilities (Annex D);
- d) Support Aggregation Port Debug Information package support as specified in 7.3 and in this case shall support the Churn Detection machine as specified in 6.4.17;
- e) Conform to the required specifications of the Link Layer Discovery Protocol (LLDP) of IEEE Std 802.1AB. A conformant System that supports LLDP shall provide an LLDP Parser/Multiplexer to attach zero or more instances of LLDP to each Aggregation Port as specified in 6.1.3;
- f) Implement the Protocol Parser/Multiplexer (6.2.7) for a protocol not explicitly specified in 6.2.7;
- g) Support Per-service frame distribution as specified in 8.2. A conformant System that supports Per-service frame distribution shall support:
 - 1) Conversation-sensitive frame collection and distribution state diagrams as specified in 6.6.1; and
 - 2) Classification by Customer VLAN Tag for C-tagged interfaces and classification by Service VLAN tag for S-tagged and B-tagged interfaces, and may support:
 - 3) Conversation-sensitive LACP operation as specified in 6.6.2;
 - 4) Classification by Backbone Service Instance Tag as specified in 8.2.2;
 - 5) Classification by other methods not specified by this standard;
- h) Transmit and receive version 2 LACPDUs in the formats specified in 6.4.2. A conformant System that supports version 2 LACPDUs shall support the Long LACPDU machine as specified in 6.4.18.

5.4 Distributed Resilient Network Interconnect requirements

A conformant System, shall:

- a) Conform to the requirements specified for Link Aggregation as specified in 5.3;
- b) Support a Portal consisting of two Portal Systems as specified in Clause 9. A conformant System that supports a Portal consisting of two Portal Systems shall:
 - 1) Support a DR Function having a single IPP and conform to the state machines and procedures in 9.2 and 9.3 that are applicable to a Portal System supporting a single IPP;
 - 2) Support the Distributed Relay Control Protocol and conform to the state machine and procedures applicable to a Portal System supporting a single IPP as specified in 9.4;
 - 3) Transmit and receive required DRCPDUs in the formats specified in 9.4.3 that are applicable to a Portal consisting of two Portal Systems;
- c) Support using separate physical links for the IPL and the network links as specified in 9.3.2.

5.4.1 Distribution Resilient Network Interconnect options

A conformant System, may implement any of the options specified for Link Aggregation (5.3.1), and may:

- a) Support a Portal consisting of three Portal Systems as specified in Clause 9. A conformant System that supports a Portal consisting of three Portal Systems shall:
 - 1) Support a DR Function having two IPPs and conform to the state machines and procedures in 9.2 and 9.3 that are applicable to a Portal System supporting two IPPs;
 - 2) Support the Distributed Relay Control Protocol and conform to the state machine and procedures applicable to a Portal System supporting two IPPs as specified in 9.4;
 - 3) Transmit and receive required DRCPDUs in the formats specified in 9.4.3 for three Portal Systems;
- b) Support any of the methods in 9.3.2, by which the Systems can distinguish frames on a network link from frames on a particular Intra-Portal Link. A System that supports Network / IPL sharing by time (9.3.2.1) may also support MAC address synchronization as specified in Annex G.

- 1 c) Transmit and receive DRCPDUs in the formats specified in item b3) in 5.4, including any additional
- 2 DRCP TLVs from Table 9-8.
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18
- 19
- 20
- 21
- 22
- 23
- 24
- 25
- 26
- 27
- 28
- 29
- 30
- 31
- 32
- 33
- 34
- 35
- 36
- 37
- 38
- 39
- 40
- 41
- 42
- 43
- 44
- 45
- 46
- 47
- 48
- 49
- 50
- 51
- 52
- 53
- 54

6. Link Aggregation

6.1 Overview

This clause defines an optional Link Aggregation sublayer for use with links offering the IEEE Std 802.1AC Internal Sublayer Service (ISS). Link Aggregation allows one or more links to be aggregated together to form a Link Aggregation Group, such that an Aggregator Client can treat the Link Aggregation Group as if it were a single link. To this end, Link Aggregation specifies the establishment of DTE to DTE logical links, consisting of N parallel instances of full duplex point-to-point links.

The models of operation in this clause provide a basis for specifying the externally observable behavior of the operation, and are not intended to place additional constraints on implementations; these can adopt any internal model of operation compatible with the externally observable behavior specified. Conformance of equipment to this standard is purely in respect of observable protocol.

6.1.1 Goals and objectives

Link Aggregation, as specified in this clause, provides the following:

- a) **Increased bandwidth**—The capacity of multiple links is combined into one logical link.
- b) **Linearly incremental bandwidth**—Bandwidth can be increased in unit multiples as opposed to the order-of-magnitude increase available through Physical Layer technology options (10 Mb/s, 100 Mb/s, 1000 Mb/s, etc.).
- c) **Increased availability**—The failure or replacement of a single link within a Link Aggregation Group need not cause failure from the perspective of an Aggregator Client.
- d) **Load sharing**—Aggregator Client traffic may be distributed across multiple links.
- e) **Automatic configuration**—In the absence of manual overrides, an appropriate set of Link Aggregation Groups is automatically configured, and individual links are allocated to those groups. If a set of links can aggregate, they will aggregate.
- f) **Rapid configuration and reconfiguration**—In the event of changes in physical connectivity, Link Aggregation will quickly converge to a new configuration, typically on the order of milliseconds for link down events and 1 second or less for link up events.

NOTE 1—Timing out LACPDU exchanges is the method of last resort for detecting link failures and shifting data flows to other physical links. The MAC_Operational status generated by the link hardware is the first choice for link failure detection. IEEE Std 802.1Q Clause 18 Connectivity Fault Management provides a method for detecting link failures (also indicated via the MAC_Operational status) when hardware means are not applicable.

- g) **Deterministic behavior**—Depending on the selection algorithm chosen, the configuration can be made to resolve deterministically; i.e., the resulting aggregation can be made independent of the order in which events occur, and be completely determined by the capabilities of the individual links and their physical connectivity.
- h) **Low risk of duplication or misordering**—During both steady-state operation and link (re)configuration, there is a low probability that frames are duplicated or misordered.
- i) **Support of existing MAC Clients**—No change is required to existing higher-layer protocols or applications to use Link Aggregation.
- j) **Backwards compatibility with aggregation-unaware devices**—Links that cannot take part in Link Aggregation—either because of their inherent capabilities, management configuration, or the capabilities of the devices to which they attach—operate as normal, individual links.
- k) **Accommodation of differing capabilities and constraints**—Devices with differing hardware and software constraints on Link Aggregation are, to the extent possible, accommodated.
- l) **No change to frame formats**—Link aggregation neither adds to, nor changes the contents of frames exchanged between Aggregator Clients.

- m) **Network management support**—The standard specifies appropriate management objects for configuration, monitoring, and control of Link Aggregation.
- n) **Dissimilar MACs**—Link Aggregation can be used over any physical or logical medium offering the ISS.

Link Aggregation, as specified in this clause, does not support the following:

- o) **Multipoint Aggregations**—The mechanisms specified in this clause do not support aggregations among more than two Systems. (See Clause 9, Distributed Resilient Network Interconnect, for aggregations among two Portals each consisting of up to three Systems.)
- p) **Half-duplex operation**—Link Aggregation is supported only on point-to-point links with MACs operating in full duplex mode.

Aggregation of links of different data rates is not prohibited nor required by this standard. Determining how to distribute traffic across links of different data rates is beyond the scope of this standard.

NOTE 2—Previous versions of this standard restricted the data rate of the aggregated links. However, the specified mechanisms did not depend on whether the links operated at the same rate or not.

6.1.2 Positioning of Link Aggregation within the IEEE 802 architecture

Link Aggregation comprises an optional sublayer between an Aggregator Client and the MAC. Figure 6-1 depicts the positioning of the Link Aggregation sublayer in the IEEE 802 layer architecture, and the relationship of that architecture to the Data Link and Physical Layers of the OSI Reference Model. The figure also shows the ability of the Link Aggregation sublayer to combine a number of individual links in order to present a single MAC interface to the Aggregator Client. (See also IEEE Std 802.1Q-2011, Figure 22-8.)

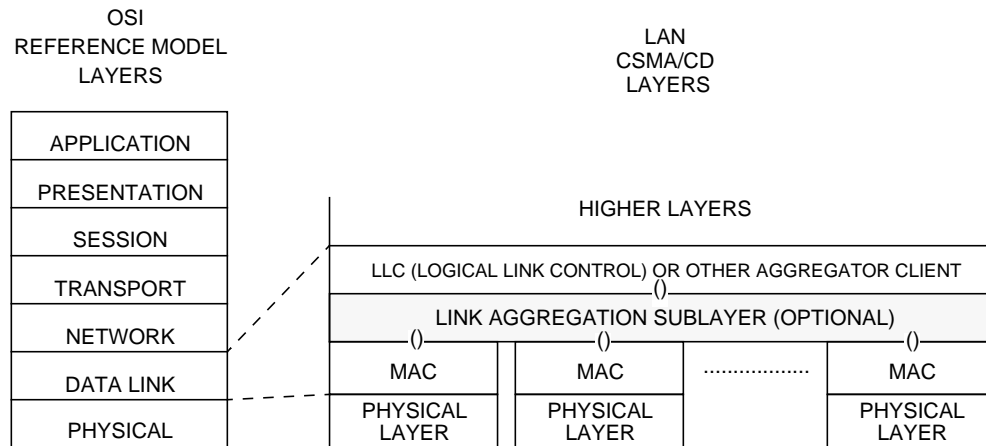


Figure 6-1—Architectural positioning of Link Aggregation sublayer

Figure 6-2 depicts the major blocks that form the Link Aggregation sublayer, and their interrelationships.

It is possible to implement the optional Link Aggregation sublayer for some ports within a System while not implementing it for other ports; i.e., it is not necessary for all ports in a System to be subject to Link Aggregation. A conformant implementation is not required to be able to apply the Link Aggregation sublayer to every port.

NOTE—Earlier editions of IEEE 802.1AX and IEEE 802.3 Clause 43 used the term “MAC Client” instead of the term “Aggregator Client,” and placed Link Aggregation between the IEEE 802.3 MAC sublayer and the IEEE 802.3 MAC Control sublayer. This Standard repositions Link Aggregation immediately above the MAC sublayer. This rearrangement of the abstract layering diagram is intended to align the Standard with common industry practice; it does not cause an implementation that is conformant to earlier editions of the Standard to become non-conformant with this edition.

6.1.3 LLDP Parser/Multiplexer

Not shown in Figure 6-1 is the fact that some protocols require that the higher layers be able to access physical ports directly, without passing through the Link Aggregation sublayer. Some protocols specify a means whereby this access is accomplished. See, for example, IEEE Std 802.1X-2010, Figure 6-2, the “SecY”.

In order to satisfy the needs of other protocols, the Protocol Parser/Multiplexer (6.2.7) is provided. In particular, if IEEE Std 802.1AB is implemented in a System supporting Link Aggregation, an LLDP Parser/Multiplexer shall be provided to attach zero or more instances of LLDP to each Aggregation Port as configured by the system administrator.

There are N LLDP Parser/Multiplexers, one for each Aggregation Port. Each LLDP Parser/Multiplexer is an instance of the Protocol Parser/Multiplexer described in 6.2.7, and is inserted as a shim layer, either between the Control Parser/Multiplexer and the MAC layer, or between the Aggregator Parser/Multiplexer and the Control Parser/Multiplexer, at the discretion of the implementer. Specifically:

- a) The *DownPort* of Protocol Parser/Multiplexer N is either the *DataMuxN* of Control Parser/Multiplexer N (6.2.10) or the *MacN* of MAC N, as determined by the position of the LLDP Parser/Multiplexer in the protocol stack.
- b) The *ControlPort* of Protocol Parser/Multiplexer N is utilized by the LLDP instances.
- c) The *DataPort* of Protocol Parser/Multiplexer N is either the *DataMuxN* of Aggregator Parser/Multiplexer N (6.2.8) or the *MacN* of Control Parser/Multiplexer N, as determined by the position of the LLDP Parser/Multiplexer in the protocol stack.
- d) The *IsControlFrame* function is defined in 6.1.3.1.1.

NOTE—Any functional entity that operates above the Aggregator is under the control and the specification of the System implementing Link Aggregation (in Bridges, for example, this is provided by the Transmit and Receive functions specified in Clause 8.5 of IEEE Std 802.1Q-2011). The discussion of the position of any Protocol Parser/Multiplexer entity that is above the Aggregator is not in scope of this standard.

6.1.3.1 LLDP Parser state diagram

6.1.3.1.1 LLDP Parser Function

IsControlFrame

Returns Boolean value: (Length/Type == LLDP_Type && DA == LLDP_Multicast address)

6.1.3.1.2 Constants

LLDP_Multicast address

The value of the LLDP DA field. (See IEEE Std 802.1AB-2009 Clause 7.1.)

LLDP_Type

The value of the LLDP Length/Type field. (See IEEE Std 802.1AB-2009 Figure C.1.)

6.1.3.1.3 Variables

DA

The value of the DA field in a received frame.

Length/Type

The value of the Length/Type field in a received frame.

6.1.3.1.4 State diagram

The LLDP Parser state diagram, when supported, shall implement the function specified in Figure 6-5 with its associated parameters as specified in 6.1.3.

6.2 Link Aggregation operation

As depicted in Figure 6-2, the Link Aggregation sublayer comprises the following functions:

- a) *Frame Distribution*. This block is responsible for taking frames submitted by the Aggregator Client and submitting them for transmission on the appropriate Aggregation Port, based on a frame distribution algorithm employed by the Frame Distributor. Frame Distribution includes a *Frame Distributor* and an optional *Marker Generator/Receiver* used for the Marker protocol. (See 6.2.4, 6.2.5, and 6.5.)
- b) *Frame Collection*. This block is responsible for passing frames received from the various Aggregation Ports to the Aggregator Client. Frame Collection includes a *Frame Collector* and a *Marker Responder*, used for the Marker protocol. (See 6.2.3, 6.2.6 and 6.5.)
- c) *Aggregator Parser/Multiplexers*. On transmit, these blocks simply pass frame transmission requests from the Frame Distributor, Marker Generator, and/or Marker Responder to the appropriate Aggregation Port. On receive, these blocks distinguish among Marker Request, Marker Response, and MAC Protocol Data Units (MPDUs), and pass each to the appropriate entity (Marker Responder, Marker Receiver, and Frame Collector, respectively).
- d) *Aggregator*. The combination of Frame Distribution and Frame Collection, along with the Aggregator Parser/Multiplexers, is referred to as the Aggregator.
- e) *Aggregation Control*. This block is responsible for the configuration and control of Link Aggregation. It incorporates a *Link Aggregation Control Protocol (LACP)* that can be used for automatic communication of aggregation capabilities between Systems and automatic configuration of Link Aggregation.
- f) *Control Parser/Multiplexers*. On transmit, these blocks simply pass frame transmission requests from the Aggregator and Control entities to the appropriate Aggregation Port. On receive, these blocks distinguish Link Aggregation Control PDUs from other frames, passing the LACPDU to the appropriate sublayer entity, and all other frames to the Aggregator.

6.2.1 Principles of Link Aggregation

Link Aggregation allows an Aggregator Client to treat a set of one or more Aggregation Ports as if it were a single port. In doing so, it employs the following principles and concepts:

- a) An Aggregator Client communicates with a set of Aggregation Ports through an Aggregator, which presents a standard ISS interface to the Aggregator Client. The Aggregator binds to one or more Aggregation Ports within a System.
- b) It is the responsibility of the Aggregator to distribute frame transmissions from the Aggregator Client to the various Aggregation Ports, and to collect received frames from the Aggregation Ports and pass them to the Aggregator Client transparently.

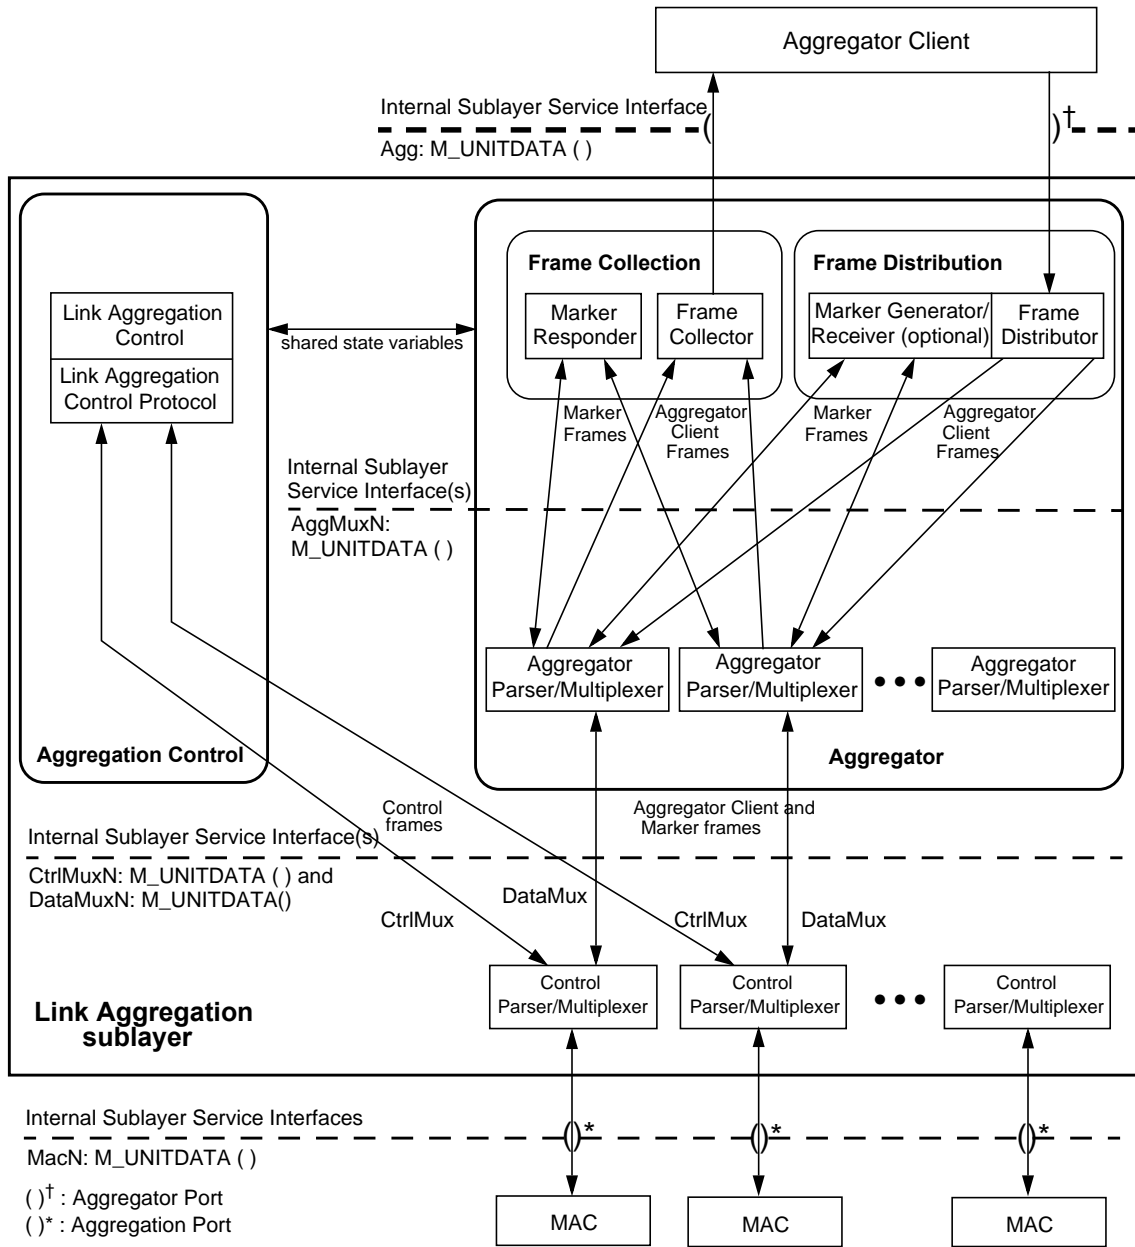


Figure 6-2—Link Aggregation sublayer block diagram

- c) A System may contain multiple Aggregators, serving multiple Aggregator Clients. A given Aggregation Port will bind to (at most) a single Aggregator at any time. An Aggregator Client is served by a single Aggregator at a time.
- d) The binding of Aggregation Ports to Aggregators within a System is managed by the Link Aggregation Control function for that System, which is responsible for determining which links may be aggregated, aggregating them, binding the Aggregation Ports within the System to an appropriate Aggregator, and monitoring conditions to determine when a change in aggregation is needed.
- e) Such determination and binding may be under manual control through direct manipulation of the state variables of Link Aggregation (e.g., Keys) by a network manager. In addition, automatic determination, configuration, binding, and monitoring may occur through the use of a Link Aggregation Control Protocol (LACP). The LACP uses peer exchanges across the links to

- determine, on an ongoing basis, the aggregation capability of the various links, and continuously provides the maximum level of aggregation capability achievable between a given pair of Systems.
- f) Frame ordering has to be maintained for certain sequences of frame exchanges between Aggregator Clients (known as conversations, see Clause 3). The Frame Distributor ensures that all frames of a given conversation are passed to a single Aggregation Port. For any given Aggregation Port, the Frame Collector is required to pass frames to the Aggregator Client in the order that they are received from that Aggregation Port. The Frame Collector is otherwise free to select frames received from the Aggregation Ports in any order. Since there are no means for frames to be misordered on a single link, this guarantees that frame ordering is maintained for any conversation.
 - g) Conversations may be moved among Aggregation Ports within an aggregation, both for load balancing and to maintain availability in the event of link failures. Means are provided (6.5, 6.6) for maintaining frame ordering when such movement takes place.
 - h) This standard does not impose any particular distribution algorithm on the Frame Distributor. Whatever algorithm is used should be appropriate for the Aggregator Client being supported.
 - i) Each Aggregation Port is assigned a MAC address, unique over the Link Aggregation Group and the 802.1Q Bridged LAN (if any) to which the Link Aggregation Group is connected. This MAC address is used as the source address for frame exchanges that are initiated by entities within the Link Aggregation sublayer itself (i.e., LACP and Marker protocol exchanges).
NOTE—The LACP and Marker protocols use a multicast destination address for all exchanges, and do not impose any requirement for an Aggregation Port to recognize more than one unicast address on received frames.¹
 - j) Each Aggregator is assigned a MAC address, unique over the Link Aggregation Group and the 802.1Q Bridged LAN (if any) to which the Link Aggregation Group is connected; this address is used as the MAC address of the aggregation from the perspective of the Aggregator Client, both as a source address for transmitted frames and as the destination address for received frames. The MAC address of the Aggregator may be one of the MAC addresses of an Aggregation Port in the associated Link Aggregation Group (see 6.2.11).

6.2.2 Service interfaces

The Aggregator Client communicates with the Aggregator using the ISS specified in IEEE Std 802.1AC. Similarly, Link Aggregation communicates internally (between Frame Collection/Distribution, the Aggregator Parser/Multiplexers, the Control Parser/Multiplexers, and Link Aggregation Control) and with its bound Aggregation Ports using the same service interface. No new interlayer service interfaces are defined for Link Aggregation.

Since Link Aggregation uses four instances of the ISS, it is necessary to introduce a notation convention so that the reader can be clear as to which interface is being referred to at any given time. A prefix is therefore assigned to each service primitive, indicating which of the four interfaces is being invoked, as depicted in Figure 6-2. The prefixes are as follows:

- a) *Agg.*, for primitives issued on the interface between the Aggregator Client and the Link Aggregation sublayer.
- b) *AggMuxN.*, for primitives issued on the interface between Aggregator Parser/Multiplexer N and its internal clients (where N is the Port Number associated with the Aggregator Parser/Multiplexer).
- c) *CtrlMuxN.*, for primitives issued on the interface between Control Parser/Multiplexer N and Link Aggregation Control (where N is the Port Number associated with the Control Parser/Multiplexer).
- d) *DataMuxN.*, for primitives issued on the interface between Control Parser/Multiplexer N and Aggregator Parser/Multiplexer N (where N is the Port Number associated with the Control Parser/Multiplexer).

1. Notes in text, tables, and figures are given for information only and do not contain requirements needed to implement the standard.

- 1 e) *MacN:*, for primitives issued on the interface between underlying MAC N and its Control Parser/
2 Multiplexer (where N is the Port Number associated with the underlying MAC).
3

4 Aggregator Clients may generate *Agg:M_UNITDATA.request* primitives for transmission on an aggregated
5 link. These are passed by the Frame Distributor to an Aggregation Port selected by the distribution
6 algorithm. *MacN:M_UNITDATA.indication* primitives signifying received frames are passed unchanged
7 from an Aggregation Port to the Aggregator Client by the Frame Collector.
8

9 In order to satisfy the needs of higher layers to access physical ports directly, the Protocol Parser/Multiplexer
10 (6.2.7) is provided. In particular, if IEEE Std 802.1AB is implemented in a System supporting Link
11 Aggregation, the LLDP Parser/Multiplexer (6.1.3) will enable to attach zero or more instances of LLDP to
12 each Aggregation Port.
13

14 6.2.3 Frame Collector

15 A Frame Collector is responsible for receiving incoming frames (i.e.,
16 *AggMuxN:M_UNITDATA.indications*) from the set of individual links that form the Link Aggregation
17 Group (through each link's associated Aggregator Parser/Multiplexer) and delivering them to the
18 Aggregator Client. Frames received from a given Aggregation Port are delivered to the Aggregator Client in
19 the order that they are received by the Frame Collector. Since the Frame Distributor is responsible for
20 maintaining any frame ordering constraints, there is no requirement for the Frame Collector to perform any
21 reordering of frames received from multiple links. If a System uses conversation-sensitive frame collection
22 and distribution then the Frame Collector will discard frames received on Aggregation Ports whose
23 Conversation IDs are not in the list supplied by the Frame Distributor for that Aggregation Port (6.6).
24
25

26 The Frame Collector shall implement the function specified in the state diagram shown in Figure 6-3 and the
27 associated definitions contained in 6.2.3.1.
28

29 6.2.3.1 Frame Collector state diagram

30 6.2.3.1.1 Constants

31 *CollectorMaxDelay*

32 In tens of microseconds, the maximum time that the Frame Collector may delay the delivery of
33 a frame received from an Aggregator Parser to its Aggregator Client. Value is assigned by
34 management or administration policy.
35
36 Value: Integer
37
38

39 6.2.3.1.2 Variables

40 *DA*

41 *SA*

42 *mac_service_data_unit*

43 *priority*

44 The parameters of the *M_UNITDATA.indication* primitive.
45

46 *BEGIN*

47 A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set
48 to FALSE when (re-)initialization has completed.
49

50 Value: Boolean
51

52 6.2.3.1.3 Messages

53 *Agg:M_UNITDATA.indication*

54 *AggMuxN:M_UNITDATA.indication*

The service primitives used to pass a received frame to a client with the specified parameters.

6.2.3.1.4 State diagram

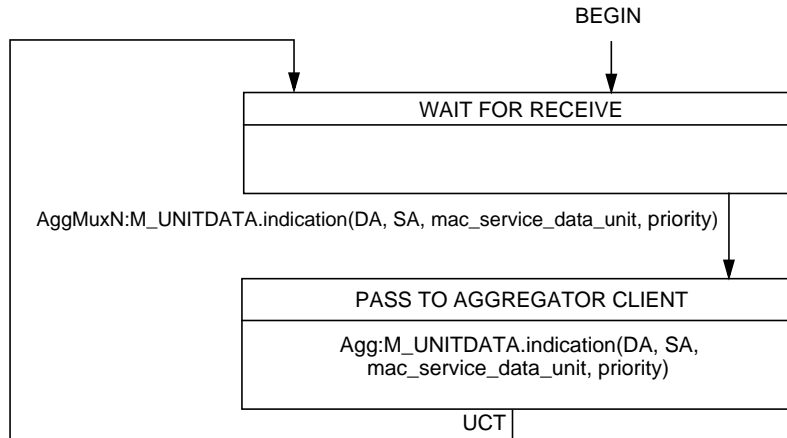


Figure 6-3—Frame Collector state diagram

The architectural models of ISS and Link Aggregation do not make any provision for queuing of frames between the link and the Aggregator Client. However, practical implementations of Link Aggregation will typically incur both queuing and delay in the Frame Collector. In order to ensure that frame delivery is not delayed indefinitely (which could cause a frame ordering problem when moving conversations from one link to another), the Frame Collector shall, upon receiving a frame from an Aggregator Parser, either deliver the frame to its Aggregator Client, or discard the frame within a CollectorMaxDelay time. The Frame Distributor (within the Partner System at the other end of the link) can assume that all frames transmitted on a given link have been either received by its Partner's Aggregator Client or discarded after a CollectorMaxDelay plus the propagation delay of the link. The use of CollectorMaxDelay is further discussed in B.3.

NOTE—Because frame discard due to CollectorMaxDelay is a function of queuing and other entities not specified in this document, the discard action is a requirement of the system, not just of the Frame Collector, and is therefore not shown in Figure 6-3.

6.2.4 Frame Distributor

The Frame Distributor is responsible for taking outgoing frames from the Aggregator Client and transmitting them through the set of links that form the Link Aggregation Group. The Frame Distributor implements a distribution function (algorithm) responsible for choosing the link to be used for the transmission of any given frame or set of frames.

This standard does not mandate any particular distribution algorithm(s); however, any distribution algorithm shall ensure that, when frames are received by a Frame Collector as specified in 6.2.3, the algorithm shall not cause

- a) Misordering of frames that are part of any given conversation, or
- b) Duplication of frames.

1 The above requirement to maintain frame ordering is met by ensuring that all frames that compose a given
2 conversation are transmitted on a single link in the order that they are generated by the Aggregator Client;
3 hence, this requirement does not involve the addition (or modification) of any information to the MAC
4 frame, nor any buffering or processing on the part of the corresponding Frame Collector in order to reorder
5 frames. This approach permits a wide variety of distribution and load balancing algorithms to be used, while
6 also ensuring interoperability between devices that adopt differing algorithms.
7

8 NOTE—The subject of distribution algorithms and maintenance of frame ordering is discussed in Clause 8 and in Annex
9 B.

10 The Frame Distributor shall implement the function specified in the state diagram shown in Figure 6-4 and
11 the associated definitions contained in 6.2.4.1.
12

13 **6.2.4.1 Frame Distributor state diagram**

14 **6.2.4.1.1 Variables**

15 DA

16 SA

17 mac_service_data_unit

18 priority

19 The parameters of the M_UNITDATA.request primitive.

20 BEGIN

21 A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set
22 to FALSE when (re-)initialization has completed.

23 Value: Boolean
24

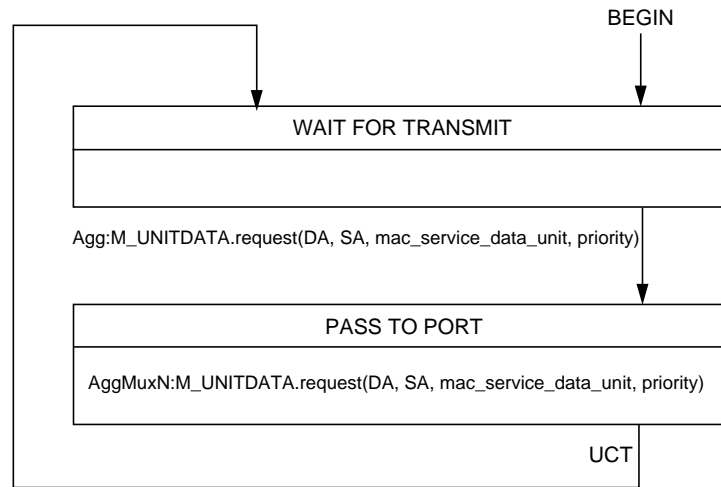
25 **6.2.4.1.2 Messages**

26 Agg:M_UNITDATA.request

27 AggMuxN:M_UNITDATA.request

28 The service primitives used to transmit a frame with the specified parameters.
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

6.2.4.1.3 State diagram



NOTE—The algorithm that the Frame Distributor uses to select the value of N in AggMuxN:M_UNITDATA.request for a given frame is unspecified.

Figure 6-4—Frame Distributor state diagram

6.2.5 Marker Generator/Receiver (optional)

The optional Marker Generator is used by the Marker protocol, as specified in 6.5. When implemented and so requested by the distribution algorithm, the Marker Generator shall issue an AggMuxN:M_UNITDATA.request primitive, with a mac_service_data_unit containing a Marker PDU as defined in 6.5.3, to the Aggregation Port associated with the conversation being marked, subject to the timing restrictions for Slow Protocols specified in IEEE Std 802.3 Annex 57A.

The optional Marker Receiver is used by the Marker protocol, as specified in 6.5. It receives Marker Response PDUs from the Aggregator Parser.

6.2.6 Marker Responder

The Marker Responder is used by the Marker protocol, as specified in 6.5. The Marker Responder receives Marker PDUs (generated by a Partner System's Marker Generator), and transmits a Marker Response PDU through the same Aggregation Port from which the Marker PDU was received. While implementation of the Marker Generator/Receiver is optional, the ability to respond to a Marker PDU (the Marker Responder) is mandatory. An implementation conformant to this clause that is not under a Distributed Relay function (DR Function, see 9.2, 9.3) shall implement the Marker Responder as specified in 6.5.4.2, thus ensuring that implementations that need to make use of the protocol can do so.

A Marker Responder that is under a DR Function is not required to respond to a Marker PDU.

NOTE—Since this standard lacks any means for coordinating Marker PDU responses among the Portal Systems comprising a Portal, it is safer to not respond to a Marker PDU, and allow the sender to time out, rather than to return a Marker PDU and run the risk of delivering frames out of order.

6.2.7 Protocol Parser/Multiplexer

A Protocol Parser/Multiplexer is a function of which there are three examples in this standard:

- a) The LLDP Parser/Multiplexer (6.1.3);
- b) The Control Parser/Multiplexer (6.2.10); and
- c) The DRCP Control Parser/Multiplexer (9.4.4).

A Protocol Parser/Multiplexer has three service interfaces, each an instance of the ISS:

- d) One *DownPort* makes use of an instance of the ISS to pass data and control frames to and from lower layers in the protocol stack.
- e) One *DataPort* offers an instance of the ISS to a higher level data function.
- f) One *ControlPort* offers an instance of the ISS to higher level control functions.

The specification of a given instance of the Protocol Parser/Multiplexer, e.g. the Control Parser/Multiplexer (6.2.10), ties these three notional service interfaces to specific service interfaces, e.g. those in 6.2.2, and provides a definition of the *IsControlFrame* function (6.2.7.1.1).

On transmission, the Protocol Multiplexer shall provide transparent pass-through of frames from the ControlPort or the DataPort to the DownPort.

On receipt of a frame from its DownPort, the Protocol Parser uses the *IsControlFrame* function to determine whether the frame is or is not a control frame. It passes control frames up through the ControlPort, and passes all other frames up through the DataPort. The Protocol Parser shall implement the function specified by the state diagram shown in Figure 6-5 and the associated definitions contained in 6.2.7.1.

6.2.7.1 Protocol Parser state diagram

6.2.7.1.1 Functions

IsControlFrame

Given an input frame (DA, SA, mac_service_data_unit, priority), returns TRUE if the frame is a control frame for the ControlPort, or FALSE, if it is a data frame for the DataPort.

Value returned: Boolean

6.2.7.1.2 Variables

DA

SA

mac_service_data_unit

priority

The parameters of the M_UNITDATA.indication primitive.

BEGIN

A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-)initialization has completed.

Value: Boolean

6.2.7.1.3 Messages

ControlPort:M_UNITDATA.indication

DataPort:M_UNITDATA.indication

DownPort:M_UNITDATA.indication

The service primitives used to pass a received frame to a client with the specified parameters.

6.2.7.1.4 State diagram

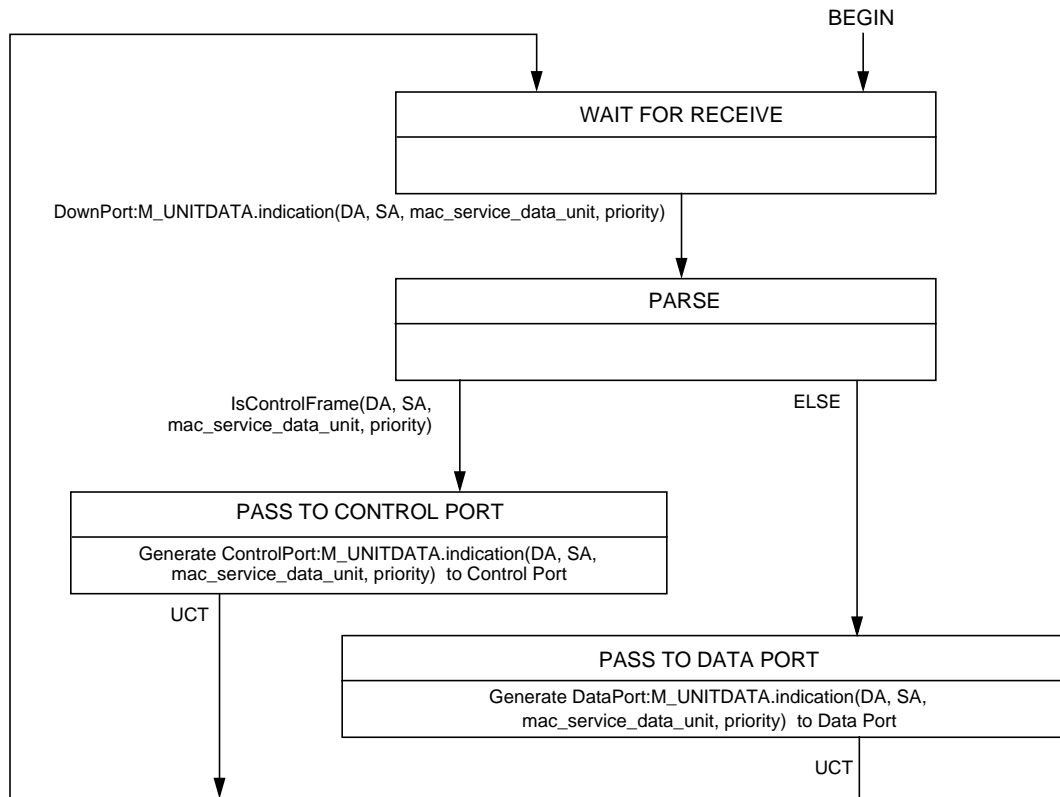


Figure 6-5—Protocol Parser state diagram

6.2.8 Aggregator Parser/Multiplexer

33
34
35 On transmission, the Aggregator Multiplexer shall provide transparent pass-through of frames submitted by
36 the Marker Responder and optional Marker Generator to the Aggregation Port specified in the transmission
37 request. The Aggregator Multiplexer shall provide transparent pass-through of frames submitted by the
38 Frame Distributor to the Aggregation Port specified in the transmission request only when the Aggregation
39 Port state is Distributing (see 6.4.15); otherwise, such frames shall be discarded.

40
41 On receipt, the Aggregator Parser decodes frames received from the Control Parser, passes those frames
42 destined for the Marker Responder or Marker Receiver to the selected entity, and discards frames with
43 invalid Slow Protocol subtype values (see IEEE Std 802.3 Table 57A-2). The Aggregator Parser shall pass
44 all other frames to the Frame Collector for passage to the Aggregator Client only when the Aggregation Port
45 state is Collecting (see 6.4.15); otherwise, such frames shall be discarded. The Aggregator Parser shall
46 implement the function specified in the state diagram shown in Figure 6-6 and the associated definitions
47 contained in 6.2.8.1.

6.2.8.1 Aggregator Parser state diagram

6.2.8.1.1 Constants

52
53 Slow_Protocols_Type

54 The value of the Slow Protocols Length/Type field. (See IEEE Std 802.3 Annex 57A.)

- 1 Marker_subtype
2 The value of the Subtype field for the Marker protocol. (See 6.5.3.)
3 Value: Integer
4 2
- 5 Marker_Information
6 The encoding of the Marker Information TLV_type field. (See 6.5.3.)
7 Value: Integer
8 1
- 9 Marker_Response_Information
10 The encoding of the Marker Response Information TLV_type field. (See 6.5.3.)
11 Value: Integer
12 2
13

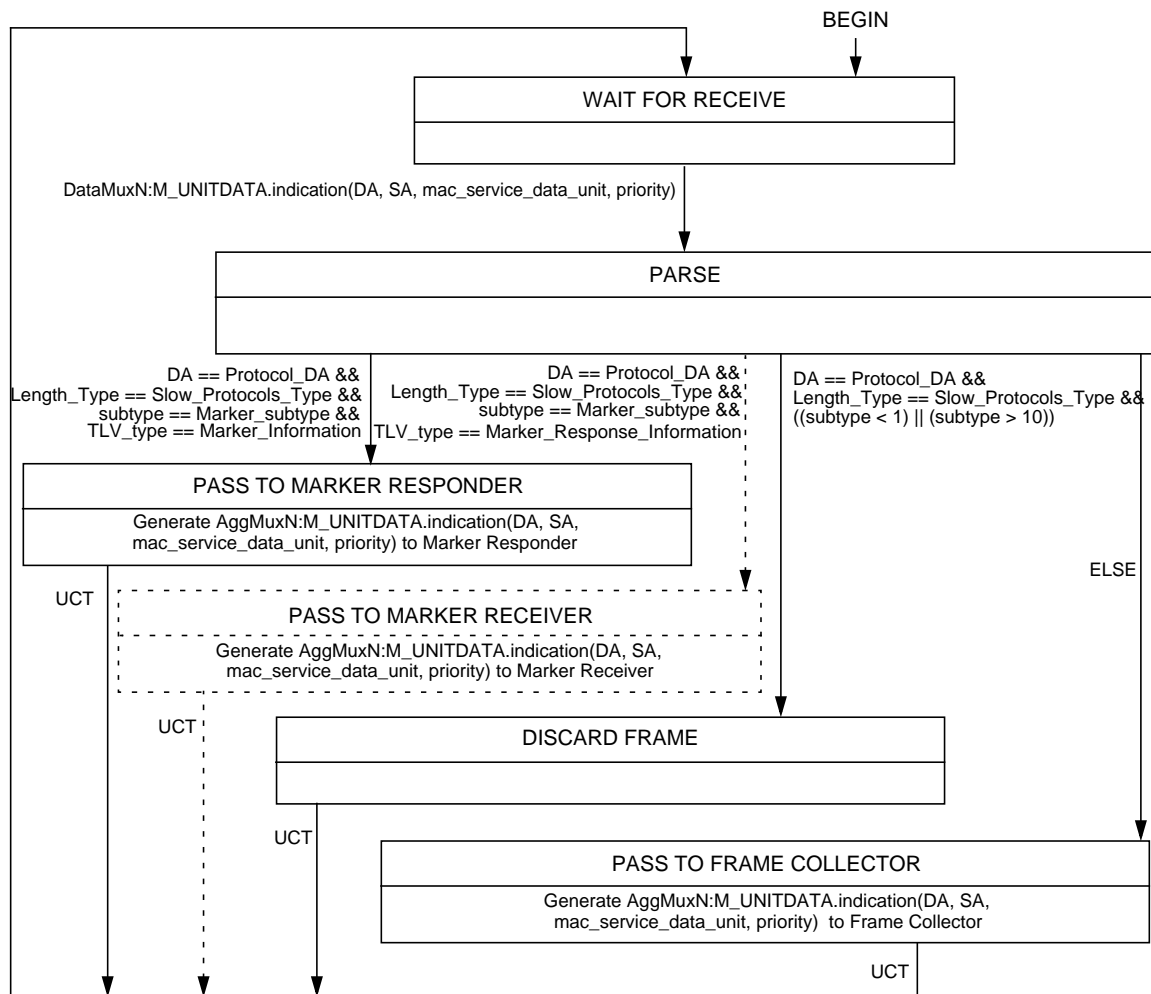
14 6.2.8.1.2 Variables

- 15
16 DA
17 SA
18 mac_service_data_unit
19 priority
20 The parameters of the M_UNITDATA.indication primitive.
21 Protocol_DA
22 One of the addresses selected from Table 6-1 determined by the setting of the
23 aAggPortProtocolDA managed object (7.3.2.2.1). A particular instance of link aggregation
24 shall use the same destination address for LACP as it does for the Marker protocol.
25 Value: 48 bits.
- 26 Length/Type
27 The value of the Length/Type field in a received frame.
28 Value: Integer
- 29 Subtype
30 The value of the octet following the Length/Type field in a Slow Protocol frame.
31 (See IEEE Std 802.3 Annex 57A.)
32 Value: Integer
- 33 TLV_type
34 The value contained in the octet following the Version Number in a received Marker or Marker
35 Response frame. This identifies the “type” for the Type/Length/Value (TLV) tuple. (See 6.5.3.)
36 Value: Integer
- 37 BEGIN
38 A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set
39 to FALSE when (re-)initialization has completed.
40 Value: Boolean
41

42 6.2.8.1.3 Messages

- 43
44 DataMuxN:M_UNITDATA.indication
45 AggMuxN:M_UNITDATA.indication
46 The service primitives used to pass a received frame to a client with the specified parameters.
47
48
49
50
51
52
53
54

6.2.8.1.4 State Diagram



If the optional Marker Receiver is not implemented, Marker Responses shall be passed to the Frame Collector. If the Aggregation Port state is not Collecting, all frames that would have been passed to the Aggregator Client through the Collector will be discarded.

Figure 6-6—Aggregator Parser state diagram

6.2.9 Aggregator

An *Aggregator* comprises an instance of a Frame Collection function, an instance of a Frame Distribution function and one or more instances of the Aggregator Parser/Multiplexer function for a Link Aggregation Group. A single Aggregator is associated with each Link Aggregation Group. An Aggregator offers an ISS interface to its associated Aggregator Client; access to the physical media by an Aggregator Client is always achieved via an Aggregator. An Aggregator can therefore be considered to be a *logical MAC*, bound to one or more Aggregation Ports, through which the Aggregator Client is provided access to the physical media.

A single, individual MAC address is associated with each Aggregator (see 6.2.11).

An Aggregator is available for use by the Aggregator Client and MAC-Operational at the Aggregator ISS will be True (6.3.12) if the following are all true:

- a) It has one or more attached Aggregation Ports.

- b) The Aggregator has not been set to a disabled state by administrative action (see 7.3.1.1.13).
- c) One or more of the attached Aggregation Ports is Collecting or Distributing (see 7.3.1.1.14).

NOTE—To simplify the modeling and description of the operation of Link Aggregation, it is assumed that there are as many Aggregators as there are Aggregation Ports in a given System; however, this is not a requirement of this standard. Aggregation of two or more Aggregation Ports consists of changing the bindings between Aggregation Ports and Aggregators such that more than one Aggregation Port is bound to a single Aggregator. The creation of any aggregations of two or more links will therefore result in one or more Aggregators that are bound to more than one Aggregation Port, and one or more Aggregators that are not bound to any Aggregation Port. An Aggregator that is not bound to any Aggregation Port appears to an Aggregator Client as a MAC interface to an inactive Aggregation Port. During times when the bindings between Aggregation Ports and Aggregators are changing, or as a consequence of particular configuration choices, there may be occasions when one or more Aggregation Ports are not bound to any Aggregator.

6.2.10 Control Parser/Multiplexer

There are N Control Parser/Multiplexers, one for each Aggregation Port. Each Control Parser/Multiplexer is an instance of the Protocol Parser/Multiplexer described in 6.2.7. Specifically:

- a) The *DownPort* of the Protocol Parser/Multiplexer is *MacN* (6.2.2) for Control Parser/Multiplexer N.
- b) The *ControlPort* of the Protocol Parser/Multiplexer is *CtrlMuxN* (6.2.2) for Control Parser/Multiplexer N.
- c) The *DataPort* of the Protocol Parser/Multiplexer is *DataMuxN* (6.2.2) for Control Parser/Multiplexer N.
- d) The *IsControlFrame* function is defined in 6.2.10.1.

6.2.10.1 Control Parser state diagram

6.2.10.1.1 Control Parser Function

IsControlFrame

Returns Boolean value: (DA == Protocol_DA && Length/Type == Slow_Protocols_Type && Subtype == LACP_subtype)

6.2.10.1.2 Constants

Slow_Protocols_Type

The value of the Slow Protocols Length/Type field. (See IEEE Std 802.3 Table 57A–2.)

LACP_subtype

The value of the Subtype field for the Link Aggregation Control Protocol. (See IEEE Std 802.3 Table 57A–3.)

Value: Integer

1

6.2.10.1.3 Variables

DA

SA

mac_service_data_unit

priority

The parameters of the *M_UNITDATA.indication* primitive.

Protocol_DA

One of the addresses selected from Table 6-1 determined by the setting of the *aAggPortProtocolDA* managed object (7.3.2.2.1). A particular instance of link aggregation shall use the same destination address for LACP as it does for the Marker protocol.

- 1 Value: 48 bits.
2 Length/Type
3 The value of the Length/Type field in a received frame.
4 Value: Integer
5 Subtype
6 The value of the octet following the Length/Type field in a Slow Protocol frame.
7 (See IEEE Std 802.3 Annex 57A.)
8 Value: Integer
9

10 6.2.11 Addressing

11
12 Each IEEE 802 MAC has an associated individual MAC address, whether that MAC is used for Link
13 Aggregation or not.

14
15 Each Aggregator to which one or more Aggregation Ports are attached has an associated individual MAC
16 address (see 6.3.3). The MAC address of the Aggregator may be the individual MAC addresses of one of the
17 MACs in the associated Link Aggregation Group, or it may be a distinct MAC address. The manner in
18 which such addresses are chosen is not otherwise constrained by this standard.
19

20 6.2.11.1 Source Address

21
22 Protocol entities sourcing frames from within the Link Aggregation sublayer (e.g., LACP and the Marker
23 protocol) use the MAC address of the MAC within an underlying Aggregation Port as the source address in
24 frames transmitted through that Aggregation Port. The Aggregator Client sees only the Aggregator and not
25 the underlying MACs, and therefore uses the Aggregator's MAC address as the source address in
26 transmitted frames. If an Aggregator Client submits a frame to the Aggregator for transmission without
27 specifying a source address, the Aggregator inserts its own MAC address as the source address for
28 transmitted frames.
29

30 NOTE—This behavior causes the Aggregator to behave the same way as a standard MAC with regard to frames
31 submitted by its client.
32

33 6.2.11.2 Destination Address

34
35 Protocol entities sourcing frames from within the Link Aggregation sublayer (e.g., LACP and the Marker
36 protocol) and the Distributed Relay control protocol entities in Clause 9 use one of the MAC addresses listed
37 in Table 6-1 as the destination address for such frames. These addresses are in the range of the C-VLAN
38 component reserved addresses (see Table 8-1 in IEEE Std 802.1Q-2011) and the S-VLAN component
39 reserved addresses (see Table 8-2 in IEEE Std 802.1Q-2011). The choice of address used determines the
40 scope of propagation of Link Aggregation control and marker PDUs within a bridged LAN, as follows:
41

- 42 a) The *Nearest Customer Bridge* group address is an address that no conformant C-VLAN component
43 or IEEE 802.1D Bridge forwards. However, this address is relayed by TPMR components and
44 S-VLAN components. Therefore, Link Aggregation control, Distributed Relay control and marker
45 PDUs received on this address represent information about stations that are not separated from the
46 recipient station by any intervening C-VLAN components or IEEE Std 802.1D Bridges. There may,
47 however, be TPMR components and/or S-VLAN components in the path between the originating
48 and receiving stations.
49

50 NOTE 1—This address makes it possible to communicate information between end stations and/or Customer Bridges
51 and for that communication to be transparent to the presence or absence of TPMRs or S-VLAN components in the
52 communication path. The scope of this address is the same as that of a customer-to-customer MACSec connection.
53
54

- b) The *Slow_Protocols_Multicast* group address is an address that no conformant Two-Port MAC Relay (TPMR) component, S-VLAN component, C-VLAN component, or IEEE 802.1D Bridge can forward. Link Aggregation control, Distributed Relay control and marker PDUs transmitted using this destination address can therefore travel no further than those stations that can be reached via a single individual LAN from the originating station. Link Aggregation control, Distributed Relay control and marker PDUs received on this address therefore represent information about stations that are attached to the same individual LAN segment as the recipient station.

NOTE 2—This address makes it possible to transmit an Link Aggregation Control, Distributed Relay Control Protocol or Marker frame containing information specific to a single individual LAN, and for that information not to be propagated further than the extent of that individual LAN.

- c) The *Nearest non-TPMR Bridge* group MAC address is an address that no conformant C-VLAN component, S-VLAN component, or IEEE 802.1D Bridge can forward. However, this address is relayed by TPMR components. Therefore, Link Aggregation control, Distributed Relay control and marker PDUs received on this address represent information about stations that are not separated from the recipient station by any intervening C-VLAN component, S-VLAN component, or IEEE 802.1D Bridge. There may, however, be one or more TPMR components in the path between the originating and receiving stations.

NOTE 3—This address makes it possible to communicate information between end stations and/or bridges and for that communication to be transparent to the presence or absence of TPMRs in the transmission path. This address is primarily intended for use within provider bridged networks.

Table 6-1—Link Aggregation protocol destination addresses

Assignment	Value
Nearest Customer Bridge group address	01-80-C2-00-00-00
IEEE 802.3 <i>Slow_Protocols_Multicast</i> group address	01-80-C2-00-00-02
Nearest non-TPMR Bridge group address	01-80-C2-00-00-03

The use and encoding of the *Slow_Protocols_Multicast* group address are specified in Annex 57A of IEEE Std 802.3-2012.

6.3 Link Aggregation Control

Link Aggregation Control configures and controls the Link Aggregation sublayer using static information local to the control function and dynamic information exchanged by means of the Link Aggregation Control Protocol.

For each Aggregation Port in the System, Link Aggregation Control

- a) Maintains configuration information (reflecting the inherent properties of the individual links as well as those established by management) to control aggregation.
- b) Exchanges configuration information with other Systems to allocate the link to a Link Aggregation Group.
- NOTE—A given link is allocated to, at most, one Link Aggregation Group at a time. The allocation mechanism attempts to maximize aggregation, subject to management controls.
- c) Attaches the Aggregation Port to the Aggregator used by the Link Aggregation Group, and detaches the Aggregation Port from the Aggregator when it is no longer used by the Group.

- 1 d) Uses information from the Partner System's Link Aggregation Control entity to enable or disable the
2 Aggregator's Frame Collector and Frame Distributor.
3

4 The operation of Link Aggregation Control involves the following activities, which are described in detail in
5 6.3.1 through 6.3.16:
6

- 7 e) Checking that candidate links can actually be aggregated.
8 f) Controlling the addition of a link to a Link Aggregation Group, and the creation of the group if
9 necessary.
10 g) Monitoring the status of aggregated links to ensure that the aggregation is still valid.
11 h) Removing a link from a Link Aggregation Group if its membership is no longer valid, and removing
12 the group if it no longer has any member links.
13

14 In order to allow Link Aggregation Control to determine whether a set of links connect to the same System,
15 and to determine whether those links are compatible from the point of view of aggregation, it is necessary to
16 be able to establish
17

- 18 i) A globally unique identifier for each System that participates in Link Aggregation (see 6.3.2).
19 j) A means of identifying the set of capabilities associated with each Aggregation Port and with each
20 Aggregator, as understood by a given System.
21 k) A means of identifying a Link Aggregation Group and its associated Aggregator.
22

23 System identification allows the detection of links that are connected in a loopback configuration (i.e., both
24 ends of the same link are connected to the same System).
25

26 **6.3.1 Characteristics of Link Aggregation Control**

27

28 Link Aggregation Control provides a configuration capability that is
29

- 30 a) **Automatic.** In the absence of manual override controls, an appropriate set of Link Aggregation
31 Groups is automatically configured, and individual links are allocated to those groups. If a set of
32 links can aggregate, they do aggregate.
33 b) **Continuous.** Manual intervention or initialization events are not a requirement for correct operation.
34 The configuration mechanism continuously monitors for changes in state that require
35 reconfiguration. The configuration functions detect and correct misconfigurations by performing
36 reconfiguration and/or by taking misconfigured links out of service.
37 c) **Deterministic.** The configuration can resolve deterministically; i.e., the configuration achieved can
38 be made independent of the order in which events occur, and be completely determined by the
39 combination of the capabilities of the individual links and their physical connectivity.
40 d) **Controllable.** The configuration capabilities accommodate devices with differing hardware and
41 software constraints on Link Aggregation.
42 e) **Compatible.** Links that cannot take part in Link Aggregation, either because of their inherent
43 capabilities or of the capabilities of the devices to which they attach, operate as normal IEEE 802
44 links. The introduction of Link Aggregation capability at one or both ends of a link should not result
45 in a degradation of the perceived performance of the link.
46 f) **Rapid.** The configuration resolves rapidly to a stable configuration. Convergence can be achieved
47 by the exchange of three LACPDUs, without dependence on timer values.
48

49 with
50

- 51 g) **Low risk of misdelivery.** The operation of the (re-)configuration functions minimizes the risk of
52 frames being delivered to the wrong Aggregator.
53 h) **Low risk of duplication or misordering.** The operation of the (re-)configuration functions
54 minimizes the risk of frame duplication and frame misordering.

- 1 i) **Low protocol overhead.** The overhead involved in external communication of configuration
2 information between devices is small.

6.3.2 System identification

6 The globally unique identifier used to identify a System shall be the concatenation of a globally
7 administered individual MAC address and the System Priority. The MAC address chosen may be the
8 individual MAC address associated with one of the Aggregation Ports of the System. In the case of DRNI
9 (Clause 9), all Portal Systems in a Portal have the same System Identifier, which is provided by the
10 concatenation of the Portal's administrated MAC address (7.4.1.1.4) and the Portal's System Priority
11 (7.4.1.1.5).

13 Where it is necessary to perform numerical comparisons between System Identifiers, each System Identifier
14 is considered to be an eight octet unsigned binary number, constructed as follows:

- 16 a) The two most significant octets of the System Identifier comprise the System Priority. The System
17 Priority value is taken to be an unsigned binary number; the most significant octet of the System
18 Priority forms the most significant octet of the System Identifier.
19 b) The third most significant octet of the System Identifier is derived from the initial octet of the MAC
20 address; the least significant bit of the octet is assigned the value of the first bit of the MAC address,
21 the next most significant bit of the octet is assigned the value of the next bit of the MAC address,
22 and so on. The fourth through eighth octets are similarly assigned the second through sixth octets of
23 the MAC address.

6.3.3 Aggregator identification

26 Each Aggregator to which one or more Aggregation Ports are attached shall be assigned a unique, globally-
27 or locally- administered individual MAC address. The MAC address assigned to the Aggregator may be the
28 same as the MAC address assigned to one of its bound Aggregation Ports. No Aggregator shall be assigned
29 a MAC address that is the same as that of an Aggregation Port bound to a different Aggregator within the
30 System. When receiving frames, an Aggregation Port is never required to recognize more than one unicast
31 address, i.e., the Aggregator's MAC address.

33 NOTE—This allows that Aggregators can be uniquely addressed, and allows (but does not require) the unique address to
34 be allocated from the same set of addresses as are assigned to the Aggregation Ports. It also acknowledges the fact that
35 locally administered addresses may be used in particular implementations or environments. The stated restriction on the
36 allocation of MAC addresses to Aggregators may have implications with regard to the choice of selection algorithm.

38 An Aggregator also shall be assigned an integer identifier that is used by Link Aggregation Control to
39 uniquely identify the Aggregator within the System. This value will typically be the same as the interface
40 identifier (ifIndex) used for management purposes.

6.3.4 Port identification

43 Link Aggregation Control uses a Port Identifier (Port ID), comprising the concatenation of a Port
44 Priority(7.3.2.1.15) and a Port Number (7.3.2.1.14), to identify the Aggregation Port. Port Numbers (and
45 hence, Port Identifiers) shall be uniquely assigned within a System. Port Number 0 shall not be assigned to
46 any Aggregation Port.

49 When it is necessary to perform numerical comparisons between Port Identifiers, each Port Identifier is
50 considered to be a four octet unsigned binary number constructed as follows:

- 51 a) The most significant and second most significant octets are the first and second most significant
52 octets of the Port Priority, respectively.

- 1 b) The third and fourth most significant octets are the first and second most significant octets of the
2 Port Number, respectively.
3

4 **6.3.5 Capability identification** 5

6 The ability of one Aggregation Port to aggregate with another is summarized by a simple integer parameter,
7 known as a Key. This facilitates communication and comparison of aggregation capabilities, which may be
8 determined by a number of factors, including
9

- 10 a) The Aggregation Port's physical characteristics, such as data rate, duplexity, and point-to-point or
11 shared medium.
12 b) Configuration constraints established by the network administrator.
13 c) Use of the Aggregation Port by higher layer protocols (e.g., assignment of Network Layer
14 addresses).
15 d) Characteristics or limitations of the Aggregation Port implementation itself.
16

17 Two Keys shall be associated with each Aggregation Port—an operational Key and an administrative Key.
18 The operational Key is the Key that is currently in active use for the purposes of forming aggregations. The
19 administrative Key allows manipulation of Key values by management. The administrative and operational
20 Keys assigned to an Aggregation Port may differ
21

- 22 e) If the operation of the implementation is such that an administrative change to a Key value cannot be
23 immediately reflected in the operational state of the Aggregation Port.
24 f) If the System supports the dynamic manipulation of Keys, as discussed in 6.7.2, either to accurately
25 reflect changes in operational capabilities of the Aggregation Port (for example, as a result of Auto-
26 Negotiation), or to provide a means of handling constraints on aggregation capability.
27

28 A given Key value is meaningful only in the context of the System that allocates it; there is no global
29 significance to Key values. Similarly, the relationship between administrative and operational Key values is
30 meaningful only in the context of the System that allocates it. When a System assigns an administrative Key
31 value to a set of Aggregation Ports, it signifies that the set of Aggregation Ports have the potential to
32 aggregate together, subject to the considerations discussed in 6.7.2. When a System assigns an operational
33 Key value to a set of Aggregation Ports, it signifies that, in the absence of other constraints, the current
34 operational state of the set of Aggregation Ports allows any subset of that set of Aggregation Ports (including
35 the entire set) to be aggregated together from the perspective of the System making the assignment. The set
36 of such Aggregation Ports that will actually be aggregated will be those that terminate at a common Partner
37 System, and for which that Partner System has assigned a common operational Key value, local to that
38 Partner. The set of Aggregation Ports in a given System that share the same operational Key value are said to
39 be members of the same Key Group.
40

41 A System may determine that a given link is not able to be aggregated with other links. Such links are
42 referred to as Individual links (as opposed to Aggregateable links). A System may declare a link to be
43 Individual if the inherent properties of the link allow its use as part of an aggregation, but the system is
44 aware of no other links that are capable of aggregating with this link (e.g., the System has allocated a unique
45 operational Key value to the link).
46

47 The capability information communicated between Systems, therefore, includes this local knowledge of the
48 aggregation capability of the link in addition to the operational Key value; i.e., whether the System considers
49 the link to be Aggregateable or Individual.
50

51 An administrative Key value and an operational Key value shall also be associated with each Aggregator.
52 The operational Key is the Key that is currently in active use for the purposes of forming aggregations. The
53 administrative Key allows manipulation of Key values by management. The values of administrative and
54 operational Key for an Aggregator may differ in the same manner as that of Aggregation Port Keys, per item

1 e) and item f) in this subclause. Aggregation Ports that are members of a given Key Group can only be
2 bound to Aggregators that share the same operational Key value.
3

4 All Keys are 16-bit identifiers. All values except the null value (all zeros) are available for local use.
5

6 NOTE—This model allows for two convenient initial configurations. The first is achieved by assigning each
7 Aggregation Port an initial administrative and operational Key value identical to its Port Number, and assigning the
8 same Port Numbers as Keys to the corresponding Aggregators for each Aggregation Port. A device with this initial
9 configuration will bring up all links as individual, non-aggregated links. The second is achieved by assigning the same
10 administrative and operational Key values to all Aggregation Ports with a common set of capabilities, and also to all
11 Aggregators. A device with this initial configuration will attempt to aggregate together any set of links that have the
12 same Partner System ID and operational Key, and for which both Systems are prepared to allow aggregation.
13

14 **6.3.6 Link Aggregation Group identification** 15

16 A Link Aggregation Group consists of either
17

- 18 a) One or more Aggregateable links that terminate in the same pair of Systems (or Portals) and whose
19 Aggregation Ports belong to the same Key Group in each System, or
- 20 b) An Individual link.
21

22 **6.3.6.1 Construction of the Link Aggregation Group Identifier** 23

24 A unique Link Aggregation Group Identifier (LAG ID) is constructed from the following parameters for
25 each of the communicating Systems:
26

- 27 a) The System Identifier
- 28 b) The operational Key assigned to the Aggregation Ports in the LAG
- 29 c) The Port Identifier, if the link is identified as an Individual link
30

31 The local System's values for these parameters shall be non-zero. In cases where the local System is unable
32 to determine the remote System's values for these parameters by exchange of protocol information,
33 administrative values are used in the construction of the LAG ID. The value of these administrative
34 parameters for the remote System may be configured as zero, provided that the Aggregation Port(s)
35 concerned are also configured to be Individual.
36

37 A compound identifier formed from the System Identifiers and Key values alone is sufficient to identify a
38 LAG comprising Aggregateable links. However, such an identifier is not sufficient for a LAG comprising a
39 single Individual link where the Partner System Identifier and operational Key may be zero. Even if these
40 are non-zero there may be multiple Individual Links with the same System Identifier and operational Key
41 combinations, and it becomes necessary to include Port Identifiers to provide unique LAG IDs.
42

43 Given that
44

- 45 d) S and T are System Identifiers,
- 46 e) K and L are the operational Keys assigned to a LAG by S and T respectively, and
- 47 f) P and Q are the Port Identifiers of the Aggregation Ports being attached if the LAG comprises a
48 single Individual Link and zero if the LAG comprises one or more Aggregateable links,
49

50 then the general form of the unique LAG ID is [(SKP), (TLQ)].
51

52 To simplify comparison of LAG IDs it is conventional to order these such that S is the numerically smaller
53 of S and T.
54

6.3.6.2 Representation of the Link Aggregation Group Identifier

In order to allow for convenient transcription and interpretation by human network personnel, this standard provides a convention for representing compound LAG IDs. Using this format

- a) All fields are written as hexadecimal numbers, two digits per octet, in canonical format.
- b) Octets are presented in order, from left to right. Within fields carrying numerical significance (e.g., priority values), the most significant octet is presented first, and the least significant octet last.
- c) Within fields that carry MAC addresses, OUIs or CIDs, successive octets are separated by dashes (-), in accordance with the hexadecimal representation for MAC addresses defined in IEEE Std 802.
- d) Parameters of the LAG ID are separated by commas.

For example, consider the parameters for the two Partners in a Link Aggregation Group shown in Table 6-2.

Table 6-2—Example Partner parameters

	Partner SKP	Partner TLQ
System Parameters (S, T)	System Priority = 0x8000 (see 6.4.2.3) System MACaddr= AC-DE-48-03-67-80	System Priority = 0x8000 (see 6.4.2.3) System MACaddr= AC-DE-48-03-FF-FF
Key Parameter (K, L)	Key = 0x0001	Key = 0x00AA
Aggregation Port Parameters (P, Q)	Port Priority = 0x80 (see 6.4.2.3) Port Number = 0x0002	Port Priority = 0x80 (see 6.4.2.3) Port Number = 0x0002

The complete LAG ID derived from this information is represented as follows, for an Individual link:

[(SKP), (TLQ)] = [(8000,AC-DE-48-03-67-80,0001,80,0002), (8000,AC-DE-48-03-FF-FF,00AA,80,0002)]

The corresponding LAG ID for a set of Aggregateable links is represented as follows:

[(SKP), (TLQ)] = [(8000,AC-DE-48-03-67-80,0001,00,0000), (8000,AC-DE-48-03-FF-FF,00AA,00,0000)]

NOTE—The difference between the two representations is that, for an Aggregateable link, the Port Identifier components are zero.

It is recommended that this format be used whenever displaying LAG ID information for use by network personnel.

6.3.7 Selecting a Link Aggregation Group

Each Aggregation Port is selected for membership in the Link Aggregation Group uniquely identified by the LAG ID (composed of operational information, both derived from local administrative parameters and received through the Link Aggregation Control Protocol). Initial determination of the LAG ID is delayed to allow receipt of such information from a peer Link Aggregation Control entity; in the event such information is not received, locally configured administrative defaults are assumed for the remote Aggregation Port's operational parameters.

Where a particular link is known to be Individual, the complete LAG ID is not required to select the Link Aggregation Group since the link will not be aggregated with any other.

6.3.8 Agreeing on a Link Aggregation Group

Before frames are distributed and collected from a link, both the local Link Aggregation Control entity and its remote peer (if present) need to agree on the Link Aggregation Group. The Link Aggregation Control Protocol allows each of the communicating entities to check their peer's current understanding of the LAG ID, and facilitates rapid exchange of operational parameters while that understanding differs from their own. The protocol entities monitor their operation and, if agreement is not reached (perhaps due to an implementation failure), management is alerted.

The ability of LACP to signal that a particular link is Individual can accelerate the use of the link since, if both Link Aggregation Control entities know that the link is Individual, full agreement on the LAG ID is not necessary.

6.3.9 Attaching a link to an Aggregator

Once a link has selected a Link Aggregation Group, Link Aggregation Control can attach that link to a compatible Aggregator. An Aggregator is compatible if

- a) The Aggregator's operational Key matches the Aggregation Port's operational Key, and
- b) All other links currently attached to the Aggregator have selected the same Link Aggregation Group.

If several compatible Aggregators exist, Link Aggregation Control may employ a locally determined algorithm, either to ensure deterministic behavior (i.e., independence from the order in which Aggregators become available) or to maximize availability of the aggregation to an Aggregator Client. If no compatible Aggregator exists, then it is not possible to enable the link until such a time as a compatible Aggregator becomes available.

NOTE—In a properly configured System, there should always be a suitable Aggregator available with the proper Key assigned to serve a newly created Link Aggregation Group, so the unavailability of a compatible Aggregator is normally a temporary state encountered while links are moved between Aggregators. However, given the flexibility of the Key scheme, and given that in some implementations there may not be enough Aggregators to service a given configuration of links, it is possible to create configurations in which there is no Aggregator available to serve a newly identified LAG, in which case the links that are members of that LAG cannot become active until such a time as the configuration is changed to free up an appropriate Aggregator.

Links that are not successful candidates for aggregation (e.g., links that are attached to other devices that cannot perform aggregation or links that have been manually configured to be non-aggregateable) are enabled to operate as individual links. For consistency of modeling, such a link is regarded as being attached to a compatible Aggregator that can only be associated with a single link. That is, from the perspective of Link Aggregation, non-aggregated links are not a special case; they compose an aggregation with a maximum membership of one link.

More than one link can select the same Link Aggregation Group within a short period of time and, as these links detach from their prior Aggregators, additional compatible Aggregators can become available. In order to avoid such events causing repeated configuration changes, Link Aggregation Control applies hysteresis to the attachment process and allows multiple links to be attached to an Aggregator at the same time.

6.3.10 Signaling readiness to transfer user data

Once a link has been attached to an Aggregator (6.3.9) compatible with the agreed-upon Link Aggregation Group (6.3.8), each Link Aggregation Control entity signals to its peer its readiness to transfer user data to and from the Aggregator's Aggregator Client. In addition to allowing time for the organization of local Aggregator resources, including the possibility that a compatible Aggregator may not exist, explicit signaling of readiness to transfer user data can be delayed to ensure preservation of frame ordering and

1 prevention of frame duplication. Link Aggregation Control will not signal readiness until it is certain that
2 there are no frames in transit on the link that were transmitted while the link was a member of a previous
3 Link Aggregation Group. This may involve the use of an explicit Marker protocol that ensures that no
4 frames remain to be received at either end of the link before reconfiguration takes place. The operation of
5 the Marker protocol is described in 6.5. The decision as to when, or if, the Marker protocol is used is entirely
6 dependent upon the nature of the distribution algorithm that is employed.
7

8 **6.3.11 Enabling the Frame Collector and Frame Distributor**

9
10 Initially, both the Frame Collector and Frame Distributor are disabled. Once the Link Aggregation Control
11 entity is ready to transfer user data using the link and its peer entity has also signaled readiness, the process
12 of enabling the link can proceed. When any Aggregation Port attached to the Frame Collection function is
13 Collecting, the Frame Collector is enabled (thus preparing it to receive frames sent over the link by the
14 remote Aggregator's Distributor) and that fact is communicated to the Partner. Once the received
15 information indicates that the remote Aggregator's Frame Collector is enabled, the Frame Distributor is also
16 enabled.
17

18 NOTE—This description assumes that the implementation is capable of controlling the state of the transmit and receive
19 functions of the MAC independently. In an implementation where this is not possible, the transmit and receive functions
20 are enabled or disabled together. The manner in which this is achieved is detailed in the description of the Mux machine
21 (see 6.4.15).
22

23 If at least one Aggregation Port's Mux in the Link Aggregation Group is Collecting, then the Receive state
24 of the corresponding Aggregator will be Enabled. If at least one Aggregation Port's Mux in the Link
25 Aggregation Group is Distributing, then the Transmit state of the corresponding Aggregator will be Enabled.
26

27 **6.3.12 MAC_Operational status**

28
29 MAC_Operational is the Boolean status in the ISS (IEEE Std 802.1AC-2012 clause 11.2) that indicates to a
30 higher layer entity that a Service Access Point is (TRUE) or is not (FALSE) available for use. MAC_Operational
31 is an indication from a physical link to Link Aggregation that the link is available, and is
32 also an indication from the Link Aggregation Sublayer to the higher layers that the sublayer is available.
33

34 MAC_Operational from physical links is an input to the port_enabled variable (6.4.7). Through this
35 variable, a link whose MAC_Operational status is FALSE is removed from a Link Aggregation Group.
36

37 The Link Aggregation Sublayer shall present its MAC_Operational status as TRUE to higher layers if and
38 only if its Receive_State and Transmit_State variables both have the value Enabled (6.4.6). The operational
39 state of the Aggregator is “up” (MAC_Operational is TRUE) if one or more of the Aggregation Ports that
40 are attached to the Aggregator are Collecting, or both Collecting and Distributing, and if the value of
41 aAggAdminState (7.3.1.1.13) for the Aggregator is also “up.” If none of the Aggregation Ports that are
42 attached to the Aggregator are Collecting and/or Distributing, or if there are no Aggregation Ports attached
43 to this Aggregator, then the operational state is “down” (MAC_Operational is FALSE). The operational state
44 of the Aggregator is reported by the aAggOperState (7.3.1.1.14) managed object.
45

46 **6.3.13 Monitoring the membership of a Link Aggregation Group**

47
48 Each link is monitored in order to confirm that the Link Aggregation Control functions at each end of the
49 link still agree on the configuration information for that link. If the monitoring process detects a change in
50 configuration that materially affects the link's membership in its current LAG, then it may be necessary to
51 remove the link from its current LAG and to move it to a new LAG.
52
53
54

6.3.14 Detaching a link from an Aggregator

An Aggregation Port may be detached from the Aggregator used by its Link Aggregation Group as a result of protocol (e.g., Key) changes, because of System constraints (e.g., exceeding a maximum allowable number of aggregated links, or device failures) at either end of the link, or because MAC_Operational of an underlying Aggregation Port is FALSE. Both classes of events will cause the LAG ID information for the link to change, and it will be necessary for Link Aggregation Control to detach the link from its current Aggregator and move it to a new LAG (if possible). At the point where the change is detected, the Collecting and Distributing states for the Aggregation Port are set to FALSE. The Frame Distribution function is informed that the link is no longer part of the group, the changed configuration information is communicated to the corresponding Link Aggregation Partner, then the Frame Collection function is informed that the link is no longer part of the group.

Once a link has been removed from its Aggregator, the link can select its new Link Aggregation Group and then attach to a compatible Aggregator, as described in 6.3.7 and 6.3.9.

Any conversation that is reallocated to a different link as a result of detaching a link from an Aggregator shall have its frame ordering preserved. This may involve the use of the Marker protocol (6.5) or other means (6.6) to ensure that no frames that form part of that conversation remain to be received at either end of the old link before the conversation can proceed on the new link.

6.3.15 Configuration and administrative control of Link Aggregation

Administrative configuration facilities allow a degree of control to be exerted over the way that links may be aggregated. In particular, administrative configuration allows

- a) Key values associated with an Aggregation Port to be identified or modified.
- b) Key values associated with an Aggregator to be identified or modified.
- c) Links to be identified as being incapable of aggregation.
- d) Link Aggregation Control Protocol parameters to be identified or modified.

6.3.16 Link Aggregation Control state information

The Link Aggregation Control function maintains the following information with respect to each link:

- a) The identifier of the Link Aggregation Group to which it currently belongs.
- b) The identifier of the Aggregator associated with that Link Aggregation Group.
- c) The status of interaction between the Frame Collection function of the Aggregator and the link (Collecting TRUE or Collecting FALSE). Collecting TRUE indicates that the receive function of this link is enabled with respect to its participation in an aggregation; i.e., received frames will be passed up to the Aggregator for collection.
- d) The status of interaction between the Frame Distribution function of the Aggregator and the link (Distributing TRUE or Distributing FALSE). Distributing TRUE indicates that the transmit function of this link is enabled with respect to its participation in an aggregation; i.e., frames may be passed down from the Aggregator's Frame Distribution function for transmission.

This state information is communicated directly between Link Aggregation Control and the Aggregator through shared state variables without the use of a formal service interface.

The Link Aggregation Control function maintains the following information with respect to each Aggregator:

- e) The status of the Frame Collection function (Receive Enabled or Receive Disabled).
- f) The status of the Frame Distribution function (Transmit Enabled or Transmit Disabled).

1 These status values are exactly the logical OR of the Collecting and Distributing status of the individual
2 links associated with that Aggregator; i.e., if one or more links in the Link Aggregation Group are
3 Collecting, then the Aggregator is Receive Enabled, and if one or more links are Distributing, then the
4 Aggregator is Transmit Enabled.
5

6 The Transmit and Receive status of the Aggregator effectively govern the point at which the Aggregator
7 becomes available for use by the Aggregator Client, or conversely, the point at which it ceases to be
8 available.
9

10 **6.4 Link Aggregation Control Protocol (LACP)**

11 The Link Aggregation Control Protocol (LACP) provides a standardized means for exchanging information
12 between Partner Systems on a link to allow their Link Aggregation Control instances to reach agreement on
13 the identity of the Link Aggregation Group to which the link belongs, move the link to that Link
14 Aggregation Group, and enable its transmission and reception functions in an orderly manner.
15
16

17 **6.4.1 LACP design elements**

18 The following considerations were taken into account during the development of the protocol described in
19 this subclause:
20
21

- 22 a) The protocol depends upon the transmission of information and state, rather than the transmission of
23 commands. LACPDUs sent by the first party (the Actor) convey to the second party (the Actor's
24 protocol Partner) what the Actor knows, both about its own state and that of the Partner.
25
- 26 b) The information conveyed in the protocol is sufficient to allow the Partner to determine what action
27 to take next.
- 28 c) Active or passive participation in LACP is controlled by LACP_Activity, an administrative control
29 associated with each Aggregation Port, that can take the value Active LACP or Passive LACP.
30 Passive LACP indicates the Aggregation Port's preference for not transmitting LACPDUs unless its
31 Partner's control value is Active LACP (i.e., a preference not to speak unless spoken to). Active
32 LACP indicates the Aggregation Port's preference to participate in the protocol regardless of the
33 Partner's control value (i.e., a preference to speak regardless).
- 34 d) Periodic transmission of LACPDUs occurs if the LACP_Activity control of either the Actor or the
35 Partner is Active LACP. These periodic transmissions will occur at either a slow or fast transmission
36 rate depending upon the expressed LACP_Timeout preference (Long Timeout or Short Timeout) of
37 the Partner System.
- 38 e) In addition to periodic LACPDU transmissions, the protocol transmits LACPDUs when there is a
39 Need To Transmit (NTT) something to the Partner; i.e., when the Actor's state changes or when it is
40 apparent from the Partner's LACPDUs that the Partner does not know the Actor's current state.
- 41 f) The protocol assumes that the rate of LACPDU loss is very low.
42

43 There is no explicit frame loss detection/retry mechanism employed by the LACP; however, if information
44 is received from the Partner indicating that it does not have up-to-date information on the Actor's state, or if
45 the next periodic transmission is due, then the Actor will transmit a LACPDU that will correctly update the
46 Partner.
47

48 **6.4.2 LACPDU structure and encoding**

49 **6.4.2.1 Transmission and representation of octets**

50 All LACPDUs comprise an integral number of octets. The bits in each octet are numbered from 0 to 7,
51 where 0 is the low-order bit. When consecutive octets are used to represent a numerical value, the most
52 significant octet is transmitted first, followed by successively less significant octets.
53
54

1 When the encoding of (an element of) a LACPDU is depicted in a diagram

- 2
- 3 a) Octets are transmitted from top to bottom.
- 4 b) Within an octet, bits are shown with bit 0 to the left and bit 7 to the right, and are transmitted from
- 5 left to right.
- 6 c) When consecutive octets are used to represent a binary number, the octet transmitted first has the
- 7 more significant value.
- 8 d) When consecutive octets are used to represent a MAC address, the least significant bit of the first
- 9 octet is assigned the value of the first bit of the MAC address, the next most significant bit the value
- 10 of the second bit of the MAC address, and so on through the eighth bit. Similarly, the least
- 11 significant through most significant bits of the second octet are assigned the value of the ninth
- 12 through seventeenth bits of the MAC address, and so on for all the octets of the MAC address.
- 13

14 **6.4.2.2 Encapsulation of LACPDUs in frames**

15

16 An LACPDU is encoded in the `mac_service_data_unit` parameter of an `M_UNITDATA.request` or

17 `M_UNITDATA.indication`. The first octets of the `mac_service_data_unit` are a protocol identifier, followed

18 by the LACPDU, followed by padding octets, if any, as required by the underlying MAC service.

19

20 Where the ISS instance used to transmit and receive frames is provided by a media access control method

21 that can support EtherType encoding directly (e.g., is an IEEE 802.3 MAC), the protocol identifier is two

22 octets in length, and the value is the Slow Protocols EtherType (hexadecimal 88-09).

23

24 Where the ISS instance is provided by a media access method that cannot directly support EtherType

25 encoding (e.g., is an IEEE 802.11 MAC), the TPID is encoded according to the rule for a Subnetwork

26 Access Protocol (Clause 10 of IEEE Std 802) that encapsulates Ethernet frames over LLC, and comprises

27 the SNAP header (hexadecimal AA-AA-03) followed by the SNAP PID (hexadecimal 00-00-00) followed

28 by the Slow Protocols EtherType (hexadecimal 88-09).

29

30 **6.4.2.3 LACPDU structure**

31

32 The LACPDU structure shall be as shown in Figure 6-7 and as further described in the following field

33 definitions:

34

- 35 a) *Subtype*. The Subtype field identifies the specific Slow Protocol being encapsulated. LACPDUs
- 36 carry the Subtype value 0x01.
- 37 b) *Version Number*. This identifies the LACP version; Version 1 implementations conformant to this
- 38 standard carry the value 0x01 and Version 2 implementations conformant to this standard carry the
- 39 value of 0x02.
- 40 c) *TLV_type = Actor Information*. This field indicates the nature of the information carried in this TLV-
- 41 tuple. Actor information is identified by the value 0x01.
- 42 d) *Actor_Information_Length*. This field indicates the length (in octets) of this TLV-tuple, Actor
- 43 information uses a length value of 20 (0x14).
- 44 e) *Actor_System_Priority*. The priority assigned to this System (by management or administration
- 45 policy), encoded as an unsigned integer.
- 46 f) *Actor_System*. The MAC address component of Actor's System ID.
- 47 g) *Actor_Key*. The operational Key value assigned to the Aggregation Port by the Actor, encoded as an
- 48 unsigned integer.
- 49 h) *Actor_Port_Priority*. The priority assigned to this Aggregation Port by the Actor (the System
- 50 sending the PDU; assigned by management or administration policy), encoded as an unsigned
- 51 integer.
- 52 i) *Actor_Port*. The Port Number assigned to the Aggregation Port by the Actor (the System sending
- 53 the PDU), encoded as an unsigned integer.
- 54

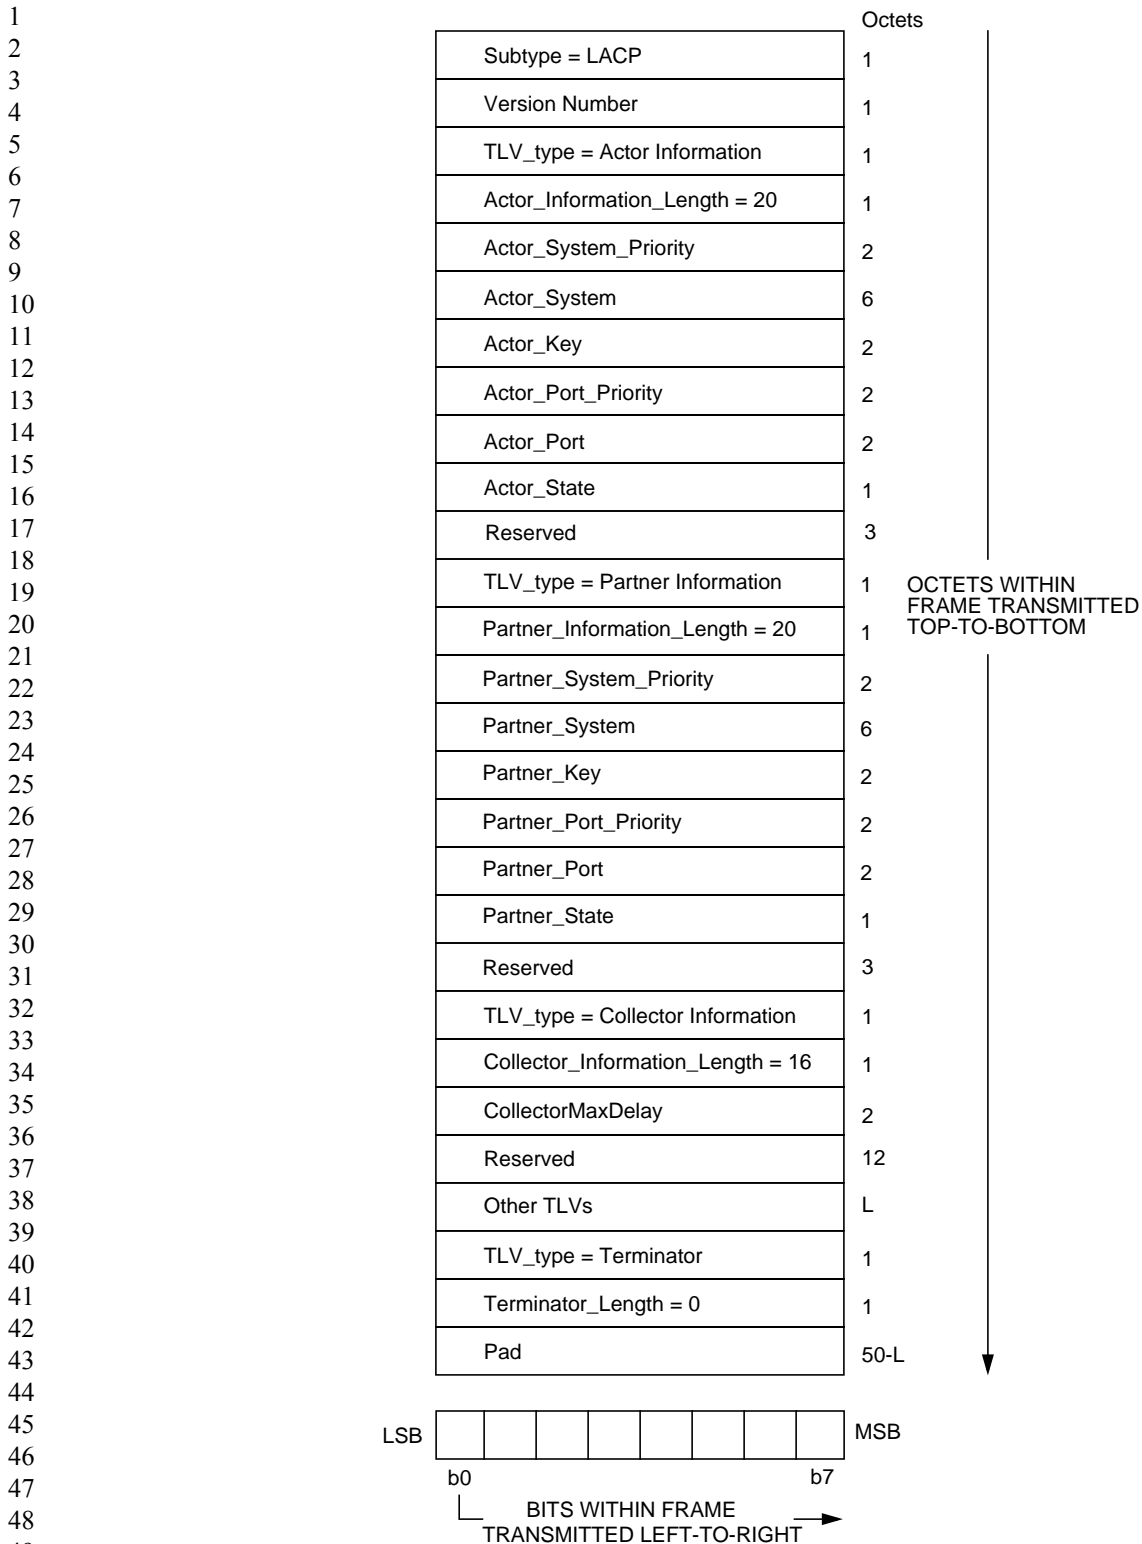


Figure 6-7—LACPDU structure

- j) *Actor_State*. The Actor’s state variables for the Aggregation Port, encoded as individual bits within a single octet, as follows and as illustrated in Figure 6-8:

- 1) *LACP_Activity* is encoded in bit 0. This flag indicates the Activity control value with regard to this link. Active LACP is encoded as a 1; Passive LACP is encoded as a 0.
- 2) *LACP_Timeout* is encoded in bit 1. This flag indicates the Timeout control value with regard to this link. Short Timeout is encoded as a 1; Long Timeout is encoded as a 0.
- 3) *Aggregation* is encoded in bit 2. If TRUE (encoded as a 1), this flag indicates that the System considers this link to be *Aggregateable*; i.e., a potential candidate for aggregation. If FALSE (encoded as a 0), the link is considered to be *Individual*; i.e., this link can be operated only as an individual link.
- 4) *Synchronization* is encoded in bit 3. If TRUE (encoded as a 1), the System considers this link to be *IN_SYNC*; i.e., it has been allocated to the correct Link Aggregation Group, the group has been associated with a compatible Aggregator, and the identity of the Link Aggregation Group is consistent with the System ID and operational Key information transmitted. If FALSE (encoded as a 0), then this link is currently *OUT_OF_SYNC*; i.e., it is not in the right Link Aggregation Group.
- 5) *Collecting* is encoded in bit 4. TRUE (encoded as a 1) means collection of incoming frames on this link is definitely enabled; i.e., collection is currently enabled and is not expected to be disabled in the absence of administrative changes or changes in received protocol information. Its value is otherwise FALSE (encoded as a 0).
- 6) *Distributing* is encoded in bit 5. FALSE (encoded as a 0) means distribution of outgoing frames on this link is definitely disabled; i.e., distribution is currently disabled and is not expected to be enabled in the absence of administrative changes or changes in received protocol information. Its value is otherwise TRUE (encoded as a 1).
- 7) *Defaulted* is encoded in bit 6. If TRUE (encoded as a 1), this flag indicates that the Actor's Receive machine is using Defaulted operational Partner information, administratively configured for the Partner. If FALSE (encoded as a 0), the operational Partner information in use has been received in a LACPDU.
- 8) *Expired* is encoded in bit 7. If TRUE (encoded as a 1), this flag indicates that the Actor's Receive machine is in the EXPIRED state; if FALSE (encoded as a 0), this flag indicates that the Actor's Receive machine is not in the EXPIRED state.

NOTE 1—The received values of Defaulted and Expired state are not used by LACP; however, knowing their values can be useful when diagnosing protocol problems.

BIT	0	1	2	3	4	5	6	7
	LACP_Activity	LACP_Timeout	Aggregation	Synchronization	Collecting	Distributing	Defaulted	Expired

NOTE 2—Bit ordering within this field is as specified in 6.4.2.1.

Figure 6-8—Bit encoding of the Actor_State and Partner_State fields

- k) *Reserved*. These 3 octets are reserved for use in future extensions to the protocol. They shall be ignored on receipt and shall be transmitted as zeros to claim compliance with Version 1 of this protocol.
- l) *TLV_type = Partner Information*. This field indicates the nature of the information carried in this TLV-tuple. Partner information is identified by the integer value 0x02.
- m) *Partner_Information_Length*. This field indicates the length (in octets) of this TLV-tuple, Partner information uses a length value of 20 (0x14).
- n) *Partner_System_Priority*. The priority assigned to the Partner System (by management or administration policy), encoded as an unsigned integer.
- o) *Partner_System*. The MAC address component of the Partner's System ID.

- 1 p) *Partner_Key*. The operational Key value assigned to the Aggregation Port associated with this link
2 by the Partner, encoded as an unsigned integer.
- 3 q) *Partner_Port_Priority*. The priority assigned to this Aggregation Port by the Partner (by
4 management or administration policy), encoded as an unsigned integer.
- 5 r) *Partner_Port*. The Port Number associated with this link assigned to the Aggregation Port by the
6 Partner, encoded as an unsigned integer.
- 7 s) *Partner_State*. The Actor's view of the Partner's state variables, depicted in Figure 6-8 and encoded
8 as individual bits within a single octet, as defined for Actor_State.
- 9 t) *Reserved*. These 3 octets are reserved for use in future extensions to the protocol. They shall be
10 ignored on receipt and shall be transmitted as zeros to claim compliance with Version 1 of this
11 protocol.
- 12 u) *TLV_type = Collector Information*. This field indicates the nature of the information carried in this
13 TLV-tuple. Collector information is identified by the integer value 0x03.
- 14 v) *Collector_Information_Length*. This field indicates the length (in octets) of this TLV-tuple.
15 Collector information uses a length value of 16 (0x10).
- 16 w) *CollectorMaxDelay*. This field contains the value of CollectorMaxDelay (6.2.3.1.1) of the station
17 transmitting the LACPDU, encoded as an unsigned integer number of tens of microseconds. The
18 range of values for this parameter is 0 to 65 535 tens of microseconds (0.65535 s).
- 19 x) *Reserved*. These 12 octets are reserved for use in future extensions to the protocol. They shall be
20 ignored on receipt and shall be transmitted as zeros to claim compliance with Version 1 of this
21 protocol.
- 22 y) *Other TLVs*. This field contains additional TLVs (Version 2 TLVs are listed in 6.4.2.4). The length *L*
23 equals the sum of the lengths of all individual additional TLVs. In Version 1 implementations of this
24 standard, *L* shall be 50 octets or less to comply with the fixed frame size of 128 octets. Version 2 and
25 later versions of the protocol do not impose such constraint and *L* can have a maximum length that is
26 only limited by the supporting MAC's maximum service data unit size (for 802.3 media, *L* can be up
27 to 1438 octets long). Version 2 or higher LACPDU's containing additional TLVs of a total length *L*
28 larger than 50 octets are Long LACPDU's.
- 29 z) *TLV_type = Terminator*. This field indicates the nature of the information carried in this TLV-tuple.
30 Terminator (end of message) information is identified by the integer value 0x00.
- 31 aa) *Terminator_Length*. This field indicates the length (in octets) of this TLV-tuple. Terminator
32 information uses a length value of 0 (0x00).
- 33 NOTE 3—The use of a Terminator_Length of 0 is intentional. In TLV encoding schemes it is common practice
34 for the terminator encoding to be 0 both for the type and the length.
35
- 36 ab) *Pad*. These 50-*L* octets are required in order to comply with the 128 octets limit imposed by Version
37 1 implementations. They are ignored on receipt and are transmitted as zeros to claim compliance
38 with Version 1 of this protocol. This field is not used in Long LACPDU's.

39 NOTE 4—Previous versions of this standard (IEEE Std 802.1AX-2008 and earlier) forced a fixed frame size of
40 128 octets on 802.3 media, regardless of the version of the protocol. This version of the standard does not
41 impose such constraint on the maximum frame size as processing limitations that were associated with earlier
42 constraints are not valid any more and the protocol is in general applicable to any media that supports the ISS.
43 In any case, a Version 1 implementation is guaranteed to be able to receive version N PDUs successfully,
44 although version N PDUs may contain additional information that cannot be interpreted (and will be ignored)
45 by the Version 1 implementation without having to force a fixed frame size. A crucial factor in ensuring
46 backwards compatibility is that any future version of the protocol is required not to redefine the structure or
47 semantics of information defined for the previous version; it may only add new information elements to the
48 previous set. Hence, in a version N PDU, a Version 1 implementation can expect to find the Version 1
49 information in exactly the same places as in a Version 1 PDU, and can expect to interpret that information as
50 defined for Version 1.
51
52
53
54

6.4.2.4 Version 2 TLVs

Table 6-3 provides a list of all TLVs that are applicable for Version 2 LACP. They support the Conversation-

Table 6-3—Type field values of Version 2 TLVs

TLV	Type Field
Port Algorithm TLV	0x04
Port Conversation ID Digest TLV	0x05
Port Conversation Mask-1	0x06
Port Conversation Mask-2	0x07
Port Conversation Mask-3	0x08
Port Conversation Mask-4	0x09
Port Conversation Service Mapping TLV	0x0A

sensitive frame collection and distribution functionality described in 6.6.

6.4.2.4.1 Port Algorithm TLV

This TLV is required to support the Conversation-sensitive LACP operation (6.6.2) and shall be present on all Version 2 LACPDUs exchanges. The Port Algorithm TLV structure shall be as shown in Figure 6-9 and as further described in the following field definitions:

TLV_type = Port Algorithm	1
Port_Algorithm_Length = 6	1
Actor_Port_Algorithm	4

Figure 6-9—Port Algorithm TLV

- TLV_type = Port Algorithm*. This field indicates the nature of the information carried in this TLV-tuple. The Port Algorithm TLV is identified by the integer value 0x04.
- Port_Algorithm_Length*. This field indicates the length (in octets) of this TLV-tuple. The Port Algorithm TLV uses a length value of 6 (0x06).
- Actor_Port_Algorithm*. This field contains the value of the algorithm used to assign frames to Port Conversation IDs. It consists of the three-octet organizationally unique identifier (OUI) or Company Identifier (CID) identifying the organization which is responsible for this algorithm and one following octet used to identify the Port Algorithm by that organization. Always set equal to aAggPortAlgorithm (7.3.1.1.33). Table 6-4 provides the IEEE 802.1 OUI (00-80-C2) Port Algorithm encodings.

Table 6-4—IEEE Port Algorithms

Port Algorithm Field	Value
Unspecified distribution algorithm	0

Table 6-4—IEEE Port Algorithms

Port Algorithm Field	Value
Distribution based on C-VIDs	1
Distribution based on S-VIDs	2
Distribution based on I-SIDs	3
Distribution based on TE-SIDs	4
Distribution based on ECMP Flow Hash	5
Reserved	6-255

6.4.2.4.2 Port Conversation ID Digest TLV

This TLV is required to support the Conversation-sensitive LACP operation (6.6.2) and shall be present on all Version 2 LACPDU exchanges. The Port Conversation ID Digest TLV structure shall be as shown in Figure 6-10 and as further described in the following field definitions:

TLV_type = Port Conversation ID Digest	1
Port_Conversation_ID_Digest_Length = 20	1
Link_Number_ID	2
Actor_Conversation_LinkList_Digest	16

Figure 6-10—Port Conversation ID Digest TLV

- a) *TLV_type = Port Conversation ID Digest*. This field indicates the nature of the information carried in this TLV-tuple. The Port Conversation ID Digest TLV is identified by the integer value 0x05.
- b) *Port_Conversation_ID_Digest_Length*. This field indicates the length (in octets) of this TLV-tuple. The Port Conversation ID Digest TLV uses a length value of 2 (0x14).
- c) *Link_Number_ID*. This field contains the operational value of the Link_Number_ID that is assigned to this Aggregation Port in order to identify Link_Number_ID configuration errors.
- d) *Actor_Conversation_LinkList_Digest*. This field contains the value of the MD5 digest Actor_Conversation_LinkList_Digest (6.6.2.1) for exchange with the Partner System.

6.4.2.4.3 Port Conversation Mask TLVs

There are four Port Conversation Mask TLVs:

- a) Port Conversation Mask-1 TLV
- b) Port Conversation Mask-2 TLV
- c) Port Conversation Mask-3 TLV
- d) Port Conversation Mask-4 TLV

If any of the Port Conversation Mask TLVs are to be carried in an LACPDU then all four shall be carried together and placed in the same order as the list above. These TLVs are required to support the Conversation-sensitive LACP operation (6.6.2) and they shall be present on all Long LACPDU exchanges.

The Port Conversation Mask-1 TLV structure shall be as shown in Figure 6-11 and as further described in the following field definitions:

TLV_type = Port Conversation Mask-1	1
Port_Conversation_Mask_1_Length = 131	1
Port_Conversation_Mask_State	1
Port_Oper_Conversation_Mask_1	128

Figure 6-11—Port Conversation Mask-1 TLV

- e) *TLV_type = Port Conversation Mask-1*. This field indicates the nature of the information carried in this TLV-tuple. The Port Conversation Mask-1 TLV is identified by the integer value 0x06.
- f) *Port_Conversation_Mask_1_Length*. This field indicates the length (in octets) of this TLV-tuple. The Port Conversation Mask-1 TLV uses a length value of 131 (0x83).
- g) *Port_Conversation_Mask_State*. The Port Conversation Mask state variables for the Aggregation Port, encoded as individual bits within a single octet, as follows and as illustrated in Figure 6-12:
- 1) *ActPar_Sync* is encoded in bit 0. This flag indicates if the Port Conversation Mask used by the Actor's Frame Distributor is the same or not as that used by the Partner's Frame Distributor. TRUE (encoded as a 1) if `Partner_Oper_Conversation_Mask == Port_Oper_Conversation_Mask`. Its value is otherwise FALSE (encoded as a 0);
 - 2) *Portal System Isolated (PSI)* is encoded in bit 1. This flag is only applicable for Portal Systems (Clause 9) and is used to indicate if the Portal System is isolated from the other Portal Systems within the Portal (9.4.8). TRUE (encoded as a 1) if `DRF_Neighbor_Oper_DRCP_State.IPP_Activity == FALSE` on all IPPs on this Portal System. Its value is otherwise FALSE (encoded as a 0);
 - 3) *Discard Wrong Conversation (DWC)* is encoded in bit 2. This flag is used to indicate if the Aggregator will discard frames with incorrect Port Conversation IDs. It is encoded as a 1 if `Discard_Wrong_Conversation == TRUE` and as 0 otherwise;
 - 4) All other bits in the octet are reserved for future use. They are transmitted as zero and are ignored on receipt.

BIT	0	1	2	3	4	5	6	7
	ActPar_Sync	PSI	DWC	Reserved	Reserved	Reserved	Reserved	Reserved

NOTE—Bit ordering within this field is as specified in 6.4.2.1.

Figure 6-12—Bit encoding of the Conversation_Mask_State fields

- h) *Port_Oper_Conversation_Mask_1*. This field contains the Boolean values of the mask for the first 1024 indexed Port Conversation IDs of the `Port_Oper_Conversation_Mask` Boolean vector encoded in increasing Port Conversation ID order starting from Port Conversation ID 0 up to Port Conversation ID 1023.

The Port Conversation Mask-2 TLV structure shall be as shown in Figure 6-13 and as further described in the following field definitions:

TLV_type = Port Conversation Mask-2	1
Port_Conversation_Mask_2_Length = 130	1
Port_Oper_Conversation_Mask_2	128

Figure 6-13—Port Conversation Mask-2 TLV

- i) *TLV_type = Port Conversation Mask-2*. This field indicates the nature of the information carried in this TLV-tuple. The Port Conversation Mask-2 TLV is identified by the integer value 0x07.
- j) *Port_Conversation_Mask_2_Length*. This field indicates the length (in octets) of this TLV-tuple. The Port Conversation Mask-2 TLV uses a length value of 130 (0x82).
- k) *Port_Oper_Conversation_Mask_2*. This field contains the Boolean values of the mask for the second 1024 indexed Port Conversation IDs of the Port_Oper_Conversation_Mask Boolean vector encoded in increasing Port Conversation ID order starting from Port Conversation ID 1024 up to Port Conversation ID 2047.

The Port Conversation Mask-3 TLV structure shall be as shown in Figure 6-14 and as further described in the following field definitions:

TLV_type = Port Conversation Mask-3	1
Port_Conversation_Mask_3_Length = 130	1
Port_Oper_Conversation_Mask_3	128

Figure 6-14—Port Conversation Mask-3 TLV

- l) *TLV_type = Port Conversation Mask-3*. This field indicates the nature of the information carried in this TLV-tuple. The Port Conversation Mask-3 TLV is identified by the integer value 0x08.
- m) *Port_Conversation_Mask_3_Length*. This field indicates the length (in octets) of this TLV-tuple. The Port Conversation Mask-3 TLV uses a length value of 130 (0x82).
- n) *Port_Oper_Conversation_Mask_3*. This field contains the Boolean values of the mask for the third 1024 indexed Port Conversation IDs of the Port_Oper_Conversation_Mask Boolean vector encoded in increasing Port Conversation ID order starting from Port Conversation ID 2048 up to Port Conversation ID 3071.

The Port Conversation Mask-4 TLV structure shall be as shown in Figure 6-15 and as further described in the following field definitions:

TLV_type = Port Conversation Mask-4	1
Port_Conversation_Mask_4_Length = 130	1
Port_Oper_Conversation_Mask_4	128

Figure 6-15—Port Conversation Mask-4 TLV

- o) *TLV_type = Port Conversation Mask-4*. This field indicates the nature of the information carried in this TLV-tuple. The Port Conversation Mask-4 TLV is identified by the integer value 0x09.
- p) *Port_Conversation_Mask_4_Length*. This field indicates the length (in octets) of this TLV-tuple. The Port Conversation Mask-4 TLV uses a length value of 130 (0x82).
- q) *Port_Oper_Conversation_Mask_4*. This field contains the Boolean values of the mask for the final 1024 indexed Port Conversation IDs of the Port_Oper_Conversation_Mask Boolean vector encoded in increasing Port Conversation ID order starting from Port Conversation ID 3072 up to Port Conversation ID 4095.

6.4.2.4.4 Port Conversation Service Mapping TLV

This TLV is only required when the Service IDs are different from the Conversation ID. The Port Conversation Service Mapping TLV structure shall be as shown in Figure 6-16 and as further described in the following field definitions:

TLV_type = Port Conversation Service Mapping Digest	1
Port_Conversation_Service_Mapping_Digest_Length = 18	1
Actor_Conversation_Service_Mapping_Digest	16

Figure 6-16—Port Conversation Service Mapping TLV

- a) *TLV_type = Port Conversation Service Mapping Digest*. This field indicates the nature of the information carried in this TLV-tuple. The Port Conversation Service Mapping TLV is identified by the integer value 0x0A.
- b) *Port_Conversation_Service_Mapping_Digest_Length*. This field indicates the length (in octets) of this TLV-tuple. The Port Conversation Service Mapping TLV uses a length value of 18 (0x12).
- c) *Actor_Conversation_Service_Mapping_Digest*. This field contains the value of the MD5 digest computed from aAggAdminServiceConversationMap[] (7.3.1.1.38) for exchange with the Partner System.

6.4.3 LACP state machine overview

The operation of the protocol is controlled by a number of state machines, each of which performs a distinct function. These state machines are for the most part described on a per-Aggregation Port basis; any deviations from per-Aggregation Port description are highlighted in the text. Events (such as expiration of a timer or received LACPDUs) may cause state transitions and also cause actions to be taken; those actions

1 may include the need for transmission of a LACPDU containing repeated or new information. Periodic and
2 event-driven transmissions are controlled by the state of a Need-To-Transmit (NTT) variable (see 6.4.7),
3 generated by the state machines as necessary.
4

5 The state machines are as follows:

- 6
- 7 a) *Receive machine (RX—6.4.12)*. This state machine receives LACPDUs from the Partner, records the
8 information contained, and times it out using either Short Timeouts or Long Timeouts, according to
9 the setting of LACP_Timeout. It evaluates the incoming information from the Partner to determine
10 whether the Actor and Partner have both agreed upon the protocol information exchanged to the
11 extent that the Aggregation Port can now be safely used, either in an aggregation with other
12 Aggregation Ports or as an individual Aggregation Port; if not, it asserts NTT in order to transmit
13 fresh protocol information to the Partner. If the protocol information from the Partner times out, the
14 Receive machine installs default parameter values for use by the other state machines.
- 15 b) *Periodic Transmission machine (6.4.13)*. This state machine determines whether the Actor and its
16 Partner will exchange LACPDUs periodically in order to maintain an aggregation (periodic
17 LACPDU exchanges occur if either or both are configured for Active LACP).
- 18 c) *Selection Logic (6.4.14)*. The Selection Logic is responsible for selecting the Aggregator to be
19 associated with this Aggregation Port.
- 20 d) *Mux machine (MUX—6.4.15)*. This state machine is responsible for attaching the Aggregation Port
21 to a selected Aggregator, detaching the Aggregation Port from a de-selected Aggregator, and for
22 turning collecting and distributing at the Aggregation Port on or off as required by the current
23 protocol information.
- 24 e) *Transmit machine (TX—6.4.16)*. This state machine handles the transmission of LACPDUs, both on
25 demand from the other state machines, and on a periodic basis.
26

27 Figure 6-17 illustrates the relationships among these state machines and the flow of information between
28 them. The set of arrows labeled Partner State Information represents new Partner information, contained in
29 an incoming LACPDU or supplied by administrative default values, being fed to each state machine by the
30 Receive machine. The set of arrows labeled Actor State Information represents the flow of updated Actor
31 state information between the state machines. Transmission of LACPDUs occurs either as a result of the
32 Periodic machine determining the need to transmit a periodic LACPDU, or as a result of changes to the
33 Actor's state information that need to be communicated to the Partner. The need to transmit a LACPDU is
34 signaled to the Transmit machine by asserting NTT. The remaining arrows represent shared variables in the
35 state machine description that allow a state machine to cause events to occur in another state machine.
36

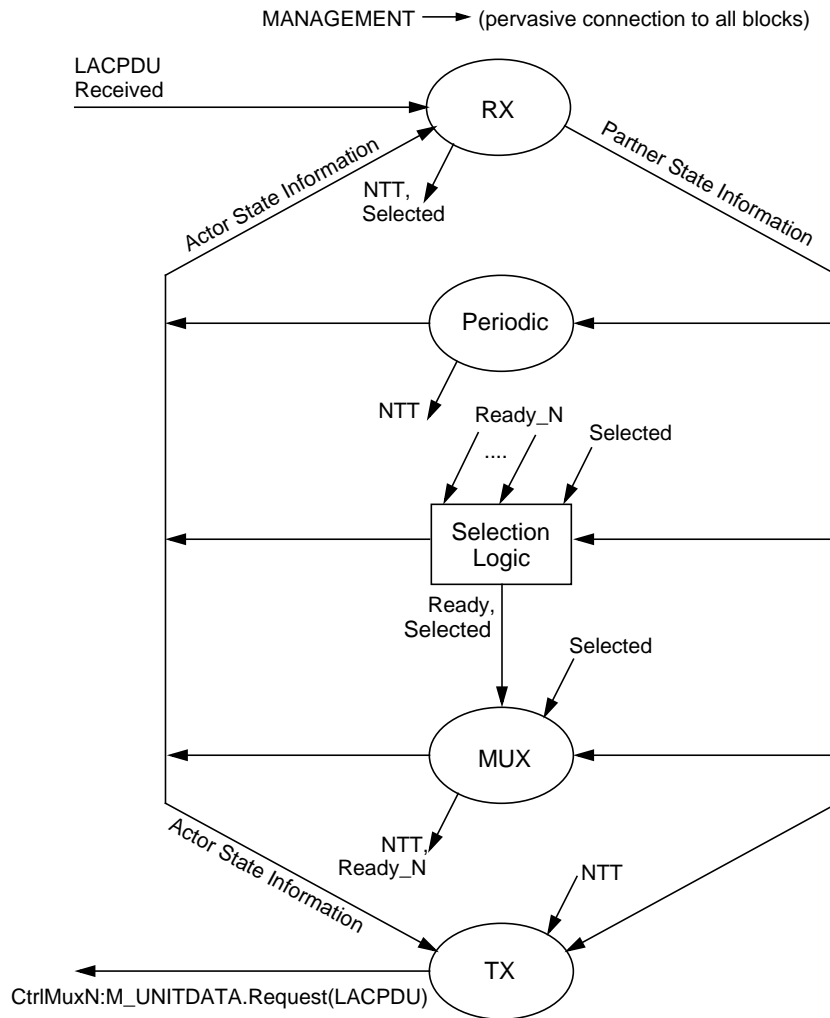
37 NOTE—The arrows marked Ready_N show that information derived from the operation of another Aggregation Port or
38 Ports can affect the operation of an Aggregation Port's state machine. See the definition of the Ready and Ready_N
39 variables in 6.4.8.
40

41 Two further state machines are defined for diagnostic purposes. They are as follows:

- 42
- 43 f) *Actor and Partner Churn Detection machines (6.4.17)*. These state machines make use of the
44 IN_SYNC and OUT_OF_SYNC states generated by the Actor and Partner Mux machines in order
45 to detect the situation where the state machines are unable to resolve the state of a given link; e.g.,
46 because the Partner System repeatedly sends conflicting information in its LACPDUs. As this
47 situation can occur in a normally functioning link, particularly where either or both participating
48 Systems have constrained aggregation capability (see 6.7), these state machines simply detect the
49 presence of such a condition and signal its existence to management.
50

51 6.4.4 Constants

52
53 All timers specified in this subclause have an implementation tolerance of ± 250 ms.
54



35 **Figure 6-17—Interrelationships among state machines**

36 **Fast_Periodic_Time**

37 The number of seconds between periodic transmissions using Short Timeouts.

38 Value: Integer

39 1

40 **Slow_Periodic_Time**

41 The number of seconds between periodic transmissions using Long Timeouts.

42 Value: Integer

43 30

44 **Short_Timeout_Time**

45 The number of seconds before invalidating received LACPDU information when using Short Timeouts ($3 \times \text{Fast_Periodic_Time}$).

46 Value: Integer

47 3

48 **Long_Timeout_Time**

49 The number of seconds before invalidating received LACPDU information when using Long Timeouts ($3 \times \text{Slow_Periodic_Time}$).

50 Value: Integer

51 90

1 Churn_Detection_Time
2 The number of seconds that the Actor and Partner Churn state machines wait for the Actor or
3 Partner Sync state to stabilize.
4 Value: Integer
5 60
6 Aggregate_Wait_Time
7 The number of seconds to delay aggregation, to allow multiple links to aggregate
8 simultaneously.
9 Value: Integer
10 2
11 Actor_System_LACP_Version
12 The Version number of the Actor's LACP implementation.
13 Value: Integer

6.4.5 Variables associated with the System

14
15
16 Actor_System
17 The MAC address component of the System Identifier of the System.
18 Value: 48 bits
19 Assigned by administrator or System policy.
20 Actor_System_Priority
21 The System Priority of the System.
22 Value: Integer
23 Assigned by administrator or System policy.
24
25

6.4.6 Variables associated with each Aggregator

26
27 Aggregator_MAC_address
28 The MAC address assigned to the Aggregator.
29 Value: 48 bits
30 Assigned by administrator or System policy.
31 Aggregator_Identifier
32 Used to uniquely identify an Aggregator within a System.
33 Value: Integer
34 Assigned by administrator or System policy.
35 Individual_Aggregator
36 The aggregation capability of the Aggregator.
37 Value: Boolean
38 TRUE if the Aggregation Port attached to this Aggregator is not capable of aggregation
39 with any other Aggregation Port.
40 FALSE if the Aggregation Port(s) attached to this Aggregator are capable of aggregation
41 with other Aggregation Ports.
42 Actor_Admin_Aggregator_Key
43 The administrative Key value associated with the Aggregator.
44 Value: Integer
45 Assigned by administrator or System policy.
46 Actor_Oper_Aggregator_Key
47 The operational Key value associated with the Aggregator.
48 Value: Integer
49 Assigned by the Actor.
50 Partner_System
51 The MAC address component of the System Identifier of the remote System to which the
52 Aggregator is connected. If the Aggregator has no attached Aggregation Ports, this variable is
53 set to 0x00-00-00-00-00-00.
54

1 Value: 48 bits
 2 Partner_System_Priority
 3 The System Priority of the remote System to which the Aggregator is connected. If the
 4 Aggregator has no attached Aggregation Ports, this variable is set to zero.
 5 Value: Integer
 6 Partner_Oper_Aggregator_Key
 7 The operational Key assigned to an aggregation by the remote System to which this Aggregator
 8 is connected. If the Aggregator has no attached Aggregation Ports, this variable is set to zero.
 9 Value: Integer
 10 Receive_State
 11 The Receive_State of the Aggregator will be Enabled if one or more Aggregation Ports
 12 attached to the Aggregator are Collecting (i.e., Actor_Oper_Port_State.Collecting is TRUE for
 13 any Aggregation Port). Otherwise, Receive_State is Disabled.
 14 Values: Enabled or Disabled
 15 Transmit_State
 16 The Transmit_State of the Aggregator will be Enabled if one or more Aggregation Ports
 17 attached to the Aggregator are Distributing (i.e., Actor_Oper_Port_State.Distributing is TRUE
 18 for any Aggregation Port). Otherwise, Transmit_State is Disabled.
 19 Values: Enabled or Disabled
 20 LAG_Ports
 21 The set of Aggregation Ports that belong to the Link Aggregation Group.
 22 Value: Integer Array
 23

24 6.4.7 Variables associated with each Aggregation Port

25
 26 Actor_Port_Number
 27 The Port Number assigned to the Aggregation Port.
 28 Value: Integer
 29 Assigned by administrator or System policy.
 30 Actor_Port_Priority
 31 The Port Priority value assigned to the Aggregation Port, used to converge dynamic Key
 32 changes.
 33 Value: Integer
 34 Assigned by administrator or System policy.
 35 Actor_Port_Aggregator_Identifier
 36 The identifier of the Aggregator to which this Aggregation Port is attached.
 37 Value: Integer
 38 NTT
 39 Need To Transmit flag.
 40 Value: Boolean
 41 TRUE indicates that there is new protocol information that should be transmitted on the
 42 link, or that the Partner needs to be reminded of the old information.
 43 FALSE otherwise.
 44 Actor_Admin_Port_Key
 45 The administrative value of Key assigned to this Aggregation Port by administrator or System
 46 policy.
 47 Value: Integer
 48 Actor_Oper_Port_Key
 49 The operational value of Key assigned to this Aggregation Port by the Actor.
 50 Value: Integer
 51 Actor_Admin_Port_State
 52 The administrative values of the Actor's state parameters. This consists of the following set of
 53 variables, as described in 6.4.2.3:
 54 LACP_Activity

1 LACP_Timeout
 2 Aggregation
 3 Synchronization
 4 Collecting
 5 Distributing
 6 Defaulted
 7 Expired
 8 Value: 8 bits
 9 Actor_Oper_Port_State
 10 The operational values of the Actor's state parameters. This consists of the following set of
 11 variables, as described in 6.4.2.3:
 12 LACP_Activity
 13 LACP_Timeout
 14 Aggregation
 15 Synchronization
 16 Collecting
 17 Distributing
 18 Defaulted
 19 Expired
 20 Value: 8 bits
 21 Partner_Admin_System
 22 Default value for the MAC address component of the System Identifier of the Partner, assigned
 23 by administrator or System policy for use when the Partner's information is unknown or
 24 expired.
 25 Value: 48 bits
 26 Partner_Oper_System
 27 The operational value of the MAC address component of the System Identifier of the Partner.
 28 The Actor sets this variable either to the value received from the Partner in an LACPDU, or to
 29 the value of Partner_Admin_System.
 30 Value: 48 bits
 31 Partner_Admin_System_Priority
 32 Default value for the System Priority component of the System Identifier of the Partner,
 33 assigned by administrator or System policy for use when the Partner's information is unknown
 34 or expired.
 35 Value: Integer
 36 Partner_Oper_System_Priority
 37 The operational value of the System Priority of the Partner. The Actor sets this variable either
 38 to the value received from the Partner in an LACPDU, or to the value of
 39 Partner_Admin_System_Priority.
 40 Value: Integer
 41 Partner_Admin_Key
 42 Default value for the Partner's Key, assigned by administrator or System policy for use when
 43 the Partner's information is unknown or expired.
 44 Value: Integer
 45 Partner_Oper_Key
 46 The operational value of the Key value assigned to this link by the Partner. The Actor sets this
 47 variable either to the value received from the Partner in an LACPDU, or to the value of
 48 Partner_Admin_Key.
 49 Value: Integer
 50 Partner_Admin_Port_Number
 51 Default value for the Port Number component of the Partner's Port Identifier, assigned by
 52 administrator or System policy for use when the Partner's information is unknown or expired.
 53 Value: Integer
 54 Partner_Oper_Port_Number

1 The operational value of the Port Number assigned to this link by the Partner. The Actor sets
2 this variable either to the value received from the Partner in an LACPDU, or to the value of
3 Partner_Admin_Port_Number.

4 Value: Integer

5 Partner_Admin_Port_Priority

6 Default value for the Port Priority component of the Partner's Port Identifier, assigned by
7 administrator or System policy for use when the Partner's information is unknown or expired.

8 Value: Integer

9 Partner_Oper_Port_Priority

10 The operational value of the priority value assigned to this link by the Partner, used to converge
11 dynamic Key changes. The Actor sets this variable either to the value received from the Partner
12 in an LACPDU, or to the value of Partner_Admin_Port_Priority.

13 Value: Integer

14 Partner_Admin_Port_State

15 Default value for the Partner's state parameters, assigned by administrator or System policy for
16 use when the Partner's information is unknown or expired. The value consists of the following
17 set of variables, as described in 6.4.2.3:

18 LACP_Activity

19 LACP_Timeout

20 Aggregation

21 Synchronization

22 Collecting

23 Distributing

24 Defaulted

25 Expired

26 The value of Collecting shall be set the same as the value of Synchronization.

27 Value: 8 bits

28 Partner_Oper_Port_State

29 The operational value of the Actor's view of the current values of the Partner's state
30 parameters. The Actor sets this variable either to the value received from the Partner in an
31 LACPDU, or to the value of Partner_Admin_Port_State. The value consists of the following set
32 of variables, as described in 6.4.2.3:

33 LACP_Activity

34 LACP_Timeout

35 Aggregation

36 Synchronization

37 Collecting

38 Distributing

39 Defaulted

40 Expired

41 Value: 8 bits

42 port_enabled

43 A variable indicating that the link has been established and the Aggregation Port is operable.

44 Value: Boolean

45 TRUE if the Aggregation Port is operable (MAC_Operational == TRUE).

46 FALSE otherwise.

47

48 NOTE—The means by which the value of the port_enabled variable is generated by the underlying MAC is
49 implementation-dependent.

50

51 Partner_LACPDU_Version_Number

52 The Version number of the LACP as reported by the Partner. The Actor sets this variable either
53 to the value received from the Partner in an LACPDU, or to the default value 1. This variable is
54 only applicable to Version 2 or higher implementations of the protocol.

1 Value: Integer
2 enable_long_pdu_xmit
3 A variable indicating that Long LACPDU's can be transmitted. TRUE when Long LACPDU's
4 can be transmitted, FALSE when only fixed sized LACPDU's (110 octets) can be transmitted.
5 The variable is only applicable to Version 2 or higher implementations of the protocol.
6 Value: Boolean
7

8 **6.4.8 Variables used for managing the operation of the state machines**

9

10 BEGIN

11 This variable indicates the initialization (or reinitialization) of the LACP protocol entity. It is
12 set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-
13)initialization has completed.

14 Value: Boolean

15 LACP_Enabled

16 This variable indicates that the Aggregation Port is operating the LACP. If the link is not a
17 point-to-point link, the value of LACP_Enabled shall be FALSE. Otherwise, the value of
18 LACP_Enabled shall be TRUE.

19 Value: Boolean
20

21 NOTE 1—Among 802.3 links, full-duplex links are always point-to-point links. Previous versions of Link Aggregation
22 (IEEE Std 802.1AX-2008, IEEE Std 802.1AXbk-2012, and IEEE Std 802.3-2005 Clause 43), which were restricted only
23 to IEEE 802.3 links, specified that the state of LACP_Enabled depended upon whether the link was operating in full
24 duplex or half duplex mode. It is the point-to-point characteristic of the link that is essential for the operation of Link
25 Aggregation, hence the above definition of LACP_Enabled.

26 actor_churn

27 This variable indicates that the Actor Churn Detection machine has detected that a local
28 Aggregation Port configuration has failed to converge within a specified time, and that
29 management intervention is required.

30 Value: Boolean

31 partner_churn

32 This variable indicates that the Partner Churn Detection machine has detected that a remote
33 Aggregation Port configuration has failed to converge within a specified time, and that
34 management intervention is required.

35 Value: Boolean

36 Ready_N

37 The MUX machine asserts Ready_N TRUE to indicate to the Selection Logic that the
38 wait_while_timer has expired and it is waiting (i.e., the Aggregation Port is in the WAITING
39 state) to attach to an Aggregator. Otherwise, its value is FALSE. There is one Ready_N value
40 for each Aggregation Port.

41 Value: Boolean

42 Ready

43 The Selection Logic asserts Ready TRUE when the values of Ready_N for all Aggregation
44 Ports that are waiting to attach to a given Aggregator are TRUE. If any of the values of
45 Ready_N for the Aggregation Ports that are waiting to attach to that Aggregator are FALSE, or
46 if there are no Aggregation Ports waiting to attach to that Aggregator, then the value of Ready
47 is FALSE.

48 Value: Boolean

49 Selected

50 A value of SELECTED indicates that the Selection Logic has selected an appropriate
51 Aggregator. A value of UNSELECTED indicates that no Aggregator is currently selected. A
52 value of STANDBY indicates that although the Selection Logic has selected an appropriate
53 Aggregator, aggregation restrictions currently prevent the Aggregation Port from being enabled
54

1 as part of the aggregation, and so the Aggregation Port is being held in a standby condition.
2 This variable can only be set to SELECTED or STANDBY by the operation of the Aggregation
3 Port's Selection Logic. It can be set to UNSELECTED by the operation of the Aggregation
4 Port's Receive machine, or by the operation of the Selection Logic associated with another
5 Aggregation Port.

6
7 NOTE 2—Setting Selected UNSELECTED in the Selection Logic associated with another Aggregation
8 Port occurs if the Selection Logic determines that the other Aggregation Port has a stronger claim to
9 attach to this Aggregation Port's current Aggregator.

10
11 Value: SELECTED, UNSELECTED, or STANDBY

12 port_moved

13 This variable is set to TRUE if the Receive machine for an Aggregation Port is in the
14 PORT_DISABLED state, and the combination of Partner_Oper_System and
15 Partner_Oper_Port_Number in use by that Aggregation Port has been received in an incoming
16 LACPDU on a different Aggregation Port. This variable is set to FALSE once the INITIALIZE
17 state of the Receive machine has set the Partner information for the Aggregation Port to
18 administrative default values.

19 Value: Boolean

20 21 6.4.9 Functions

22 recordPDU

23
24 This function records the parameter values for the Actor carried in a received LACPDU
25 (Actor_Port_Number, Actor_Port_Priority, Actor_System, Actor_System_Priority, Actor_Key,
26 and Actor_State variables) as the current Partner operational parameter values
27 (Partner_Oper_Port_Number, Partner_Oper_Port_Priority, Partner_Oper_System,
28 Partner_Oper_System_Priority, Partner_Oper_Key, and Partner_Oper_Port_State variables
29 with the exception of Synchronization) and sets Actor_Oper_Port_State.Defaulted to FALSE.

30
31 This function also updates the value of the Partner_Oper_Port_State.Synchronization using the
32 parameter values carried in received LACPDUs. Parameter values for the Partner carried in the
33 received PDU (Partner_Port, Partner_Port_Priority, Partner_System, Partner_System_Priority,
34 Partner_Key, and Partner_State.Aggregation) are compared to the corresponding operational
35 parameter values for the Actor (Actor_Port_Number, Actor_Port_Priority, Actor_System,
36 Actor_System_Priority, Actor_Oper_Port_Key, and Actor_Oper_Port_State.Aggregation).
37 Partner_Oper_Port_State.Synchronization is set to TRUE if all of these parameters match,
38 Actor_State.Synchronization in the received PDU is set to TRUE, and LACP will actively
39 maintain the link in the aggregation.

40
41 Partner_Oper_Port_State.Synchronization is also set to TRUE if the value of
42 Actor_State.Aggregation in the received PDU is set to FALSE (i.e., indicates an Individual
43 link), Actor_State.Synchronization in the received PDU is set to TRUE, and LACP will
44 actively maintain the link.

45
46 Otherwise, Partner_Oper_Port_State.Synchronization is set to FALSE.

47
48 LACP is considered to be actively maintaining the link if either the PDU's
49 Actor_State.LACP_Activity variable is TRUE or both the Actor's
50 Actor_Oper_Port_State.LACP_Activity and the PDU's Partner_State.LACP_Activity
51 variables are TRUE.

52 recordDefault

53 This function records the default parameter values for the Partner carried in the Partner Admin
54 parameters (Partner_Admin_Port_Number, Partner_Admin_Port_Priority,

1 Partner_Admin_System, Partner_Admin_System_Priority, Partner_Admin_Key, and
2 Partner_Admin_Port_State) as the current Partner operational parameter values
3 (Partner_Oper_Port_Number, Partner_Oper_Port_Priority, Partner_Oper_System,
4 Partner_Oper_System_Priority, Partner_Oper_Key, and Partner_Oper_Port_State) and sets
5 Actor_Oper_Port_State.Defaulted and Partner_Oper_Port_State.Synchronization to TRUE.

6 update_Selected

7 This function updates the value of the Selected variable, using parameter values from a newly
8 received LACPDU. The parameter values for the Actor carried in the received PDU
9 (Actor_Port, Actor_Port_Priority, Actor_System, Actor_System_Priority, Actor_Key, and
10 Actor_State.Aggregation) are compared with the corresponding operational parameter values
11 for the Aggregation Port's Partner (Partner_Oper_Port_Number, Partner_Oper_Port_Priority,
12 Partner_Oper_System, Partner_Oper_System_Priority, Partner_Oper_Key, and
13 Partner_Oper_Port_State.Aggregation). If one or more of the comparisons show that the
14 value(s) received in the PDU differ from the current operational values, then Selected is set to
15 UNSELECTED. Otherwise, Selected remains unchanged.

16 update_Default_Selected

17 This function updates the value of the Selected variable, using the Partner administrative
18 parameter values. The administrative values (Partner_Admin_Port_Number,
19 Partner_Admin_Port_Priority, Partner_Admin_System, Partner_Admin_System_Priority,
20 Partner_Admin_Key, and Partner_Admin_Port_State.Aggregation) are compared with the
21 corresponding operational parameter values for the Partner (Partner_Oper_Port_Number,
22 Partner_Oper_Port_Priority, Partner_Oper_System, Partner_Oper_System_Priority,
23 Partner_Oper_Key, and Partner_Oper_Port_State.Aggregation). If one or more of the
24 comparisons shows that the administrative value(s) differ from the current operational values,
25 then Selected is set to UNSELECTED. Otherwise, Selected remains unchanged.

26 update_NTT

27 This function updates the value of the NTT variable, using parameter values from a newly
28 received LACPDU. The parameter values for the Partner carried in the received PDU
29 (Partner_Port, Partner_Port_Priority, Partner_System, Partner_System_Priority, Partner_Key,
30 Partner_State.LACP_Activity, Partner_State.LACP_Timeout, Partner_State.Synchronization,
31 and Partner_State.Aggregation) are compared with the corresponding operational parameter
32 values for the Actor (Actor_Port_Number, Actor_Port_Priority, Actor_System,
33 Actor_System_Priority, Actor_Oper_Port_Key, Actor_Oper_Port_State.LACP_Activity,
34 Actor_Oper_Port_State.LACP_Timeout, Actor_Oper_Port_State.Synchronization, and
35 Actor_Oper_Port_State.Aggregation). If one or more of the comparisons show that the value(s)
36 received in the PDU differ from the current operational values, then NTT is set to TRUE.
37 Otherwise, NTT remains unchanged.

38 Attach_Mux_To_Aggregator

39 This function causes the Aggregation Port's Control Parser/Multiplexer to be attached to the
40 Aggregator Parser/Multiplexer of the selected Aggregator, in preparation for collecting and
41 distributing frames.

42 Detach_Mux_From_Aggregator

43 This function causes the Aggregation Port's Control Parser/Multiplexer to be detached from the
44 Aggregator Parser/Multiplexer of the Aggregator to which the Aggregation Port is currently
45 attached.

46 Enable_Collecting

47 This function causes the Aggregator Parser of the Aggregator to which the Aggregation Port is
48 attached to start collecting frames from the Aggregation Port.

49 Disable_Collecting

50 This function causes the Aggregator Parser of the Aggregator to which the Aggregation Port is
51 attached to stop collecting frames from the Aggregation Port.

52 Enable_Distributing

53 This function causes the Aggregator Multiplexer of the Aggregator to which the Aggregation
54 Port is attached to start distributing frames to the Aggregation Port.

Disable_Distributing

This function causes the Aggregator Multiplexer of the Aggregator to which the Aggregation Port is attached to stop distributing frames to the Aggregation Port.

Enable_Collecting_Distributing

This function causes the Aggregator Parser of the Aggregator to which the Aggregation Port is attached to start collecting frames from the Aggregation Port, and the Aggregator Multiplexer to start distributing frames to the Aggregation Port.

Disable_Collecting_Distributing

This function causes the Aggregator Parser of the Aggregator to which the Aggregation Port is attached to stop collecting frames from the Aggregation Port, and the Aggregator Multiplexer to stop distributing frames to the Aggregation Port.

recordVersionNumber

This function records the Partner's LACP implementation version as carried in the received LACPDU's *Version Number* field and updates the value of the variable *Partner_LACPDU_Version_Number* with the received value. This function is only applicable to Version 2 or higher implementations of this protocol.

6.4.10 Timers**current_while_timer**

This timer is used to detect whether received protocol information has expired. If *Actor_Oper_State.LACP_Timeout* is set to Short Timeout, the timer is started with the value *Short_Timeout_Time*. Otherwise, it is started with the value *Long_Timeout_Time* (see 6.4.4).

actor_churn_timer

This timer is used to detect Actor churn states. It is started using the value *Churn_Detection_Time* (see 6.4.4).

periodic_timer (time_value)

This timer is used to generate periodic transmissions. It is started using the value *Slow_Periodic_Time* or *Fast_Periodic_Time* (see 6.4.4), as specified in the Periodic Transmission state machine.

partner_churn_timer

This timer is used to detect Partner churn states. It is started using the value *Churn_Detection_Time* (see 6.4.4).

wait_while_timer

This timer provides hysteresis before performing an aggregation change, to allow all links that will join this Link Aggregation Group to do so. It is started using the value *Aggregate_Wait_Time* (see 6.4.4).

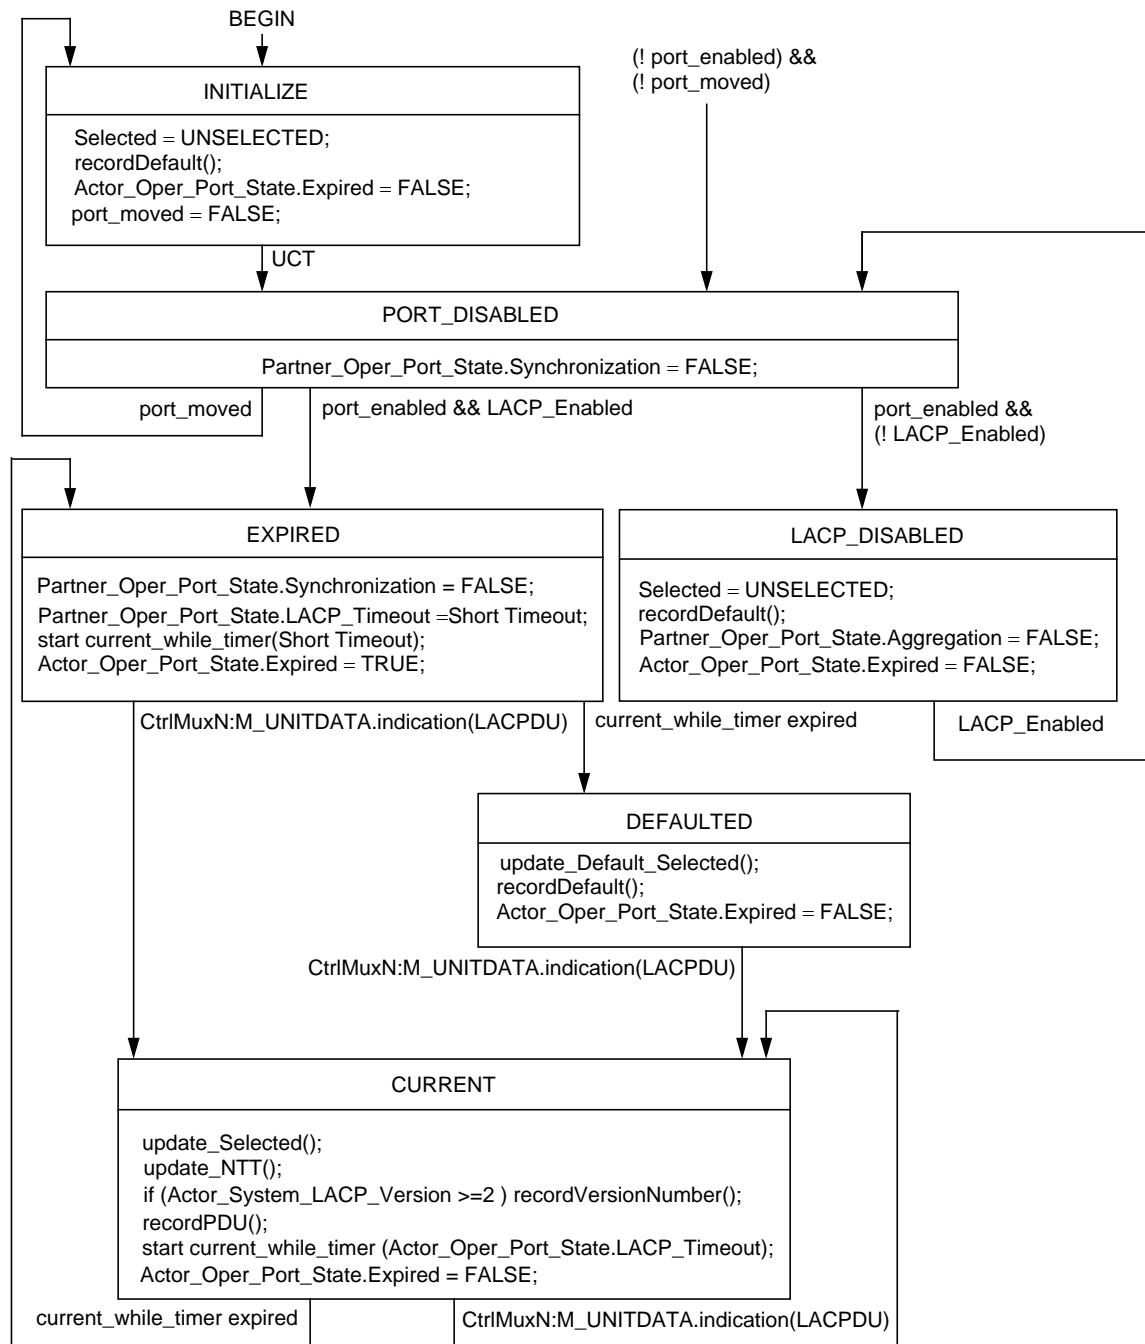
6.4.11 Messages**CtrlMuxN:M_UNITDATA.indication(LACPDU)**

This message is generated by the Control Parser as a result of the reception of a LACPDU, formatted as defined in 6.4.2.

6.4.12 Receive machine

The Receive machine shall implement the function specified in Figure 6-18 with its associated parameters (6.4.4 through 6.4.11).

On receipt of a LACPDU, the state machine enters the CURRENT state. The *update_Selected* function sets the Selected variable to UNSELECTED if the Actor's view of the Partner's operational parameters is not up to date. The Selected variable is used by the Mux machine (6.4.15).



45
46
47

Figure 6-18—Receive machine state diagram

48 NOTE 1—The Receive machine may set the Selected variable to UNSELECTED; however, setting this variable to
49 SELECTED or STANDBY is the responsibility of the Selection Logic.

50
51 The update_NTT function is used to determine whether further protocol transmissions are required; NTT is
52 set to TRUE if the Partner's view of the Actor's operational parameters is not up to date. The recordPDU
53 function records the information contained in the LACPDU in the Partner operational variables, and the
54

1 current_while timer is started. The value used to start the timer is either Short_Timeout_Time or
2 Long_Timeout_Time, depending upon the Actor's operational value of LACP_Timeout.

3
4 In the process of executing the recordPDU function, a Receive machine compliant to this standard shall not
5 validate the Version Number, TLV_type, or Reserved fields in received LACPDUs. The same actions are
6 taken regardless of the values received in these fields. A Receive machine may validate the
7 Actor_Information_Length, Partner_Information_Length, Collector_Information_Length, or
8 Terminator_Length fields. These behaviors, together with the constraint on future protocol enhancements,
9 are discussed in 6.4.2.3. Version 2 or higher implementations record as well the Version Number of the
10 received LACPs in order to check if Long LACPDUs can be transmitted on this link (6.4.18).

11
12 NOTE 2—The rules expressed above allow Version 1 devices to be compatible with future revisions of the protocol.

13
14 If no LACPDU is received before the current_while timer expires, the state machine transits to the
15 EXPIRED state. The Partner_Oper_Port_State.Synchronization variable is set to FALSE, the current
16 operational value of the Partner's LACP_Timeout variable is set to Short Timeout, and the current_while
17 timer is started with a value of Short_Timeout_Time. This is a transient state; the LACP_Timeout settings
18 allow the Actor to transmit LACPDUs rapidly in an attempt to re-establish communication with the Partner.

19
20 If no LACPDU is received before the current_while timer expires again, the state machine transits to the
21 DEFAULTED state. The update_Default_Selected function sets the Selected variable to UNSELECTED if
22 the Link Aggregation Group has changed. The recordDefault function overwrites the current operational
23 parameters for the Partner with administratively configured values. This allows configuration of
24 aggregations and individual links when no protocol Partner is present, while still permitting an active
25 Partner to override default settings. Since all operational parameters are now set to locally administered
26 values, there can be no disagreement as to the Link Aggregation Group, so the
27 Partner_Oper_Port_State.Synchronization variable is set to TRUE.

28
29 If the Aggregation Port becomes inoperable and the BEGIN variable is not asserted, the state machine enters
30 the PORT_DISABLED state. Partner_Oper_Port_State.Synchronization is set to FALSE. This state allows
31 the current Selection state to remain undisturbed, so that, in the event that the Aggregation Port is still
32 connected to the same Partner and Partner Aggregation Port when it becomes operable again, there will be
33 no disturbance caused to higher layers by unnecessary re-configuration. If the same Actor System ID and
34 Aggregation Port are seen in a LACPDU received on a different Aggregation Port (port_moved is set to
35 TRUE), this indicates that the physical connectivity has changed, and causes the state machine to enter the
36 INITIALIZE state. This state is also entered if a BEGIN event occurs.

37
38 The INITIALIZE state causes the administrative values of the Partner parameters to be used as the current
39 operational values, and sets Selected to UNSELECTED. These actions force the Mux machine to detach the
40 Aggregation Port from its current Aggregator. The variable port_moved is set to FALSE; if the entry to
41 INITIALIZE occurred as a result of port_moved being set to TRUE, then the state machine will immediately
42 transition back to the PORT_DISABLED state.

43
44 If the Aggregation Port is attached to a link that is not a point-to-point link, the operation of LACP is
45 disabled on the Aggregation Port (LACP_Enabled is FALSE) and the LACP_DISABLED state is entered.
46 This state is entered following a BEGIN or Port Enabled event. This state is similar to the DEFAULTED
47 state, except that the Aggregation Port is forced to operate as an Individual Aggregation Port, as the value of
48 Partner_Oper_Port_State.Aggregation is forced to Individual. Exit from this state occurs on a BEGIN, Port
49 Disabled, or LACP_Enabled event.

50 51 **6.4.13 Periodic Transmission machine**

52
53 The Periodic Transmission machine shall implement the function specified in Figure 6-19 with its associated
54 parameters (6.4.4 through 6.4.11).

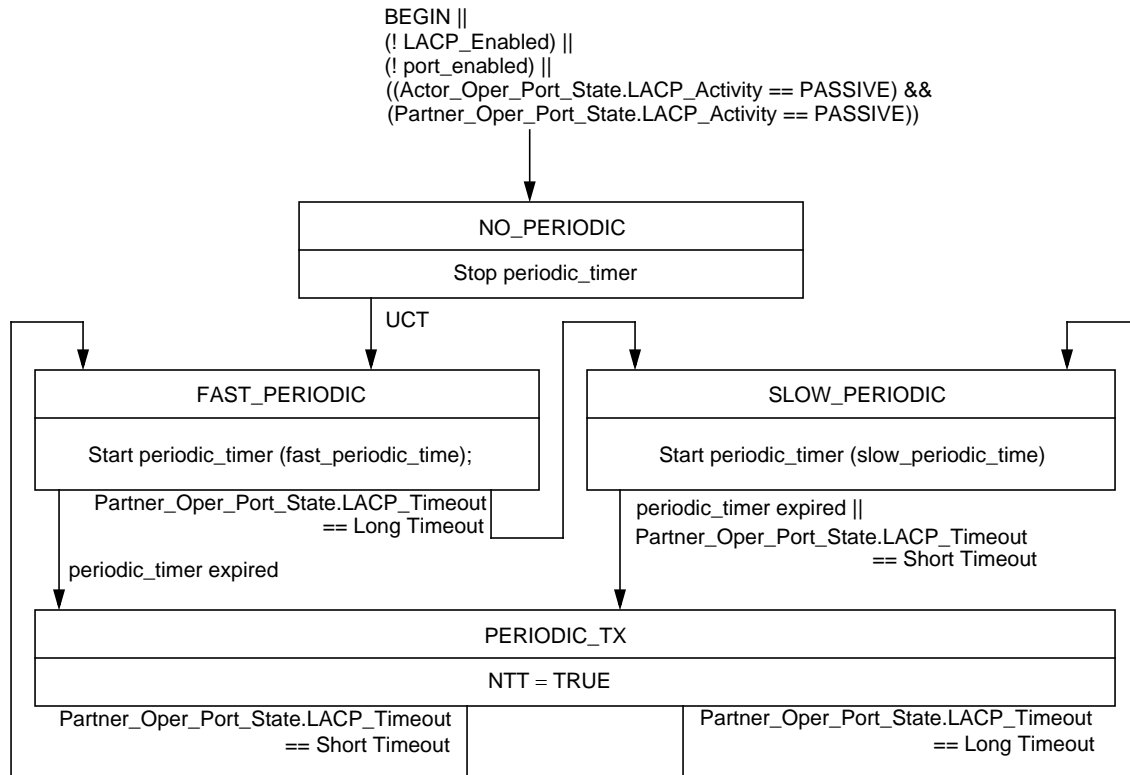


Figure 6-19—Periodic Transmission machine state diagram

The Periodic Transmission machine establishes the desire of the Actor and Partner to exchange periodic LACPDUs on the link in order to maintain an aggregation, and establishes how often those periodic transmissions should occur. Periodic transmissions will take place if either participant so wishes. Transmissions occur at a rate determined by the Partner; this rate is linked to the speed at which the Partner will time out received information.

The state machine has four states. They are as follows:

- NO_PERIODIC*. While in this state, periodic transmissions are disabled.
- FAST_PERIODIC*. While in this state, periodic transmissions are enabled at a fast transmission rate.
- SLOW_PERIODIC*. While in this state, periodic transmissions are enabled at a slow transmission rate.
- PERIODIC_TX*. This is a transitory state entered on periodic_timer expiry, that asserts NTT and then exits to *FAST_PERIODIC* or *SLOW_PERIODIC* depending upon the Partner's LACP_Timeout setting.

The values of Partner_Oper_Port_State.LACP_Activity and Actor_Oper_Port_State.LACP_Activity determine whether periodic transmissions take place. If either or both parameters are set to Active LACP, then periodic transmissions occur; if both are set to Passive LACP, then periodic transmissions do not occur. Similarly, if either of the LACP_Enabled or port_enabled variables is set to FALSE, indicating that LACP has been disabled on the Aggregation Port or that the Aggregation Port is non-operational, then no periodic transmissions take place.

If periodic transmissions are enabled, the rate at which they take place is determined by the value of the Partner_Oper_Port_State.LACP_Timeout variable. If this variable is set to Short Timeout, then the value

1 fast_periodic_time is used to determine the time interval between periodic transmissions. Otherwise,
2 slow_periodic_time is used to determine the time interval.

3 4 **6.4.14 Selection Logic**

5
6 The Selection Logic selects a compatible Aggregator for an Aggregation Port, using the Aggregation Port's
7 LAG ID. The Selection Logic may determine that the link should be operated as a standby link if there are
8 constraints on the simultaneous attachment of Aggregation Ports that have selected the same Aggregator.

9
10 NOTE 1—There will never be more than one Aggregator with the same LAG ID, but there may be none. Normally, the
11 latter will be a temporary state, caused by the fact that it takes a finite time for Aggregation Ports to be moved to the
12 correct Aggregators during reconfiguration.

13
14 The Mux machine controls the process of attaching the Aggregation Port to a selected Aggregator, after first
15 detaching the Aggregation Port from any prior Aggregator if the Aggregation Port's LAG ID has changed.

16
17 NOTE 2—An Aggregation Port is always detached from its prior Aggregator when the LAG ID changes, even if the
18 same Aggregator is selected later; to do otherwise would be to risk misdelivery of frames. Selection of a new Aggregator
19 cannot take place until the Aggregation Port is detached from any prior Aggregator; other Aggregators may become free
20 while the Aggregation Port is detaching, and other Aggregation Ports may attach to some of the available Aggregators
21 during this time interval.

22
23 The operation of the Selection Logic is separated into the following two subclauses:

- 24
25 a) The requirements for the correct operation of the Selection Logic are defined in 6.4.14.1.
26 b) The recommended default operation of the Selection Logic is described in 6.4.14.2.

27
28 This separation reflects the fact that a wide choice of selection rules is possible within the proper operation
29 of the protocol. An implementation that claims conformance to this standard may support selection rules
30 other than the recommended default; however, any such rules shall meet the requirements stated in 6.4.14.1.

31 32 **6.4.14.1 Selection Logic—Requirements**

33
34 Aggregation is represented by an Aggregation Port selecting an appropriate Aggregator, and then attaching
35 to that Aggregator. The following are required for correct operation of the selection and attachment logic:

- 36
37 a) The implementation shall support at least one Aggregator per System.
38 b) Each Aggregation Port shall be assigned an operational Key (6.3.5). Aggregation Ports that can
39 aggregate together are assigned the same operational Key as the other Aggregation Ports with which
40 they can aggregate; Aggregation Ports that cannot aggregate with any other Aggregation Port are
41 allocated unique operational Keys.
42 c) Each Aggregator shall be assigned an operational Key.
43 d) Each Aggregator shall be assigned an identifier that distinguishes it among the set of Aggregators in
44 the System.
45 e) An Aggregation Port shall only select an Aggregator that has the same operational Key assignment
46 as its own operational Key.
47 f) Subject to the exception stated in item g), Aggregation Ports that are members of the same Link
48 Aggregation Group (i.e., two or more Aggregation Ports that have the same Actor System ID, Actor
49 Key, Partner System ID, and Partner Key, and that are not required to be Individual) shall select the
50 same Aggregator.
51 g) Any pair of Aggregation Ports that are members of the same Link Aggregation Group, but are
52 connected together by the same link, shall not select the same Aggregator (i.e., if a loopback
53 condition exists between two Aggregation Ports, they shall not be aggregated together. For both
54 Aggregation Ports, the Actor System ID is the same as the Partner System ID; also, for Aggregation

Port A, the Partner's Port Identifier is Aggregation Port B, and for Aggregation Port B, the Partner's Port Identifier is Aggregation Port A).

NOTE 1—This exception condition prevents the formation of an aggregated link, comprising two ends of the same link aggregated together, in which all frames transmitted through an Aggregator are immediately received through the same Aggregator. However, it permits the aggregation of multiple links that are in loopback; for example, if Aggregation Port A is looped back to Aggregation Port C and Aggregation Port B is looped back to Aggregation Port D, then it is permissible for A and B (or A and D) to aggregate together, and for C and D (or B and C) to aggregate together.

- h) Any Aggregation Port that is required to be Individual (i.e., the operational state for the Actor or the Partner indicates that the Aggregation Port is Individual) shall not select the same Aggregator as any other Aggregation Port.
- i) Any Aggregation Port that is Aggregateable shall not select an Aggregator to which an Individual Aggregation Port is already attached.
- j) If the above conditions result in a given Aggregation Port being unable to select an Aggregator, then that Aggregation Port shall not be attached to any Aggregator.
- k) If there are further constraints on the attachment of Aggregation Ports that have selected an Aggregator, those Aggregation Ports may be selected as standby in accordance with the rules specified in 6.7.1. Selection or deselection of that Aggregator can cause the Selection Logic to re-evaluate the Aggregation Ports to be selected as standby.
- l) The Selection Logic operates upon the operational information recorded by the Receive state machine, along with knowledge of the Actor's own operational configuration and state. The Selection Logic uses the LAG ID for the Aggregation Port, determined from these operational parameters, to locate the correct Aggregator to which to attach the Aggregation Port.
- m) The Selection Logic is invoked whenever an Aggregation Port is not attached to and has not selected an Aggregator, and executes continuously until it has determined the correct Aggregator for the Aggregation Port.

NOTE 2—The Selection Logic may take a significant time to complete its determination of the correct Aggregator, as a suitable Aggregator may not be immediately available, due to configuration restrictions or the time taken to re-allocate Aggregation Ports to other Aggregators.

- n) Once the correct Aggregator has been determined, the variable Selected shall be set to SELECTED or to STANDBY (6.4.8, 6.7.1).

NOTE 3—If Selected is SELECTED, the Mux machine will start the process of attaching the Aggregation Port to the selected Aggregator. If Selected is STANDBY, the Mux machine holds the Aggregation Port in the WAITING state, ready to be attached to its Aggregator once its Selected state changes to SELECTED.

- o) The Selection Logic is responsible for computing the value of the Ready variable from the values of the Ready_N variable(s) associated with the set of Aggregation Ports that are waiting to attach to the same Aggregator (see 6.4.8).
- p) Where the selection of a new Aggregator by an Aggregation Port, as a result of changes to the selection parameters, results in other Aggregation Ports in the System being required to re-select their Aggregators in turn, this is achieved by setting Selected to UNSELECTED for those other Aggregation Ports that are required to re-select their Aggregators.

NOTE 4—The value of Selected is set to UNSELECTED by the Receive machine for the Aggregation Port when a change of LAG ID is detected.

- q) An Aggregation Port shall not be enabled for use by the Aggregator Client until it has both selected and attached to an Aggregator.
- r) An Aggregation Port shall not select an Aggregator, which has been assigned to a Portal (Clause 9), unless the Partner_Oper_Key of the associated LAG ID is equal to the lowest numerical value of the set comprising the values of the DRF_Home_Oper_Partner_Aggregator_Key, the DRF_Neighbor_Oper_Partner_Aggregator_Key and the DRF_Other_Neighbor_Oper_Partner_Aggregator_Key, on each IPP on the Portal System, where any variable having its most significant bits set to 00 is excluded (corresponding to accepting only variables that have an integer value which is larger than 16383).

- s) An Aggregation Port shall not select an Aggregator, which has been assigned to a Portal (Clause 9), if its Portal's System Identifier is set to a value that is numerically lower than the Partner's System Identifier, $PSI == TRUE$, the most significant two bits of Partner_Oper_Key are equal to the value 2 or 3 and the two least significant bits of the Aggregation Port's Partner_Oper_Port_Priority are equal to the value 2 or 3. This is to prevent network partition due to isolation of the Portal Systems in the interconnected Portals (Clause 9).

6.4.14.2 Selection Logic—Recommended default operation

The recommended default behavior provides an element of determinism (i.e., history independence) in the assignment of Aggregation Ports to Aggregators. It also has the characteristic that no additional MAC addresses are needed, over and above those already assigned to the set of underlying MACs.

NOTE—This standard does not specify any alternative selection rules beyond the recommended set. A wide variety of selection rules are possible within the scope of the requirements stated in 6.4.14.1. In particular, it is possible within these requirements to support implementations that provide fewer Aggregators than Aggregation Ports, as well as implementations designed to minimize configuration changes at the expense of less deterministic behavior.

Each Aggregation Port has an Aggregator associated with it, (i.e., the number of Aggregators in the System equals the number of Aggregation Ports supported). Each Aggregation Port/Aggregator pair is assigned the same operational Key and Port Number. When there are multiple Aggregation Ports in an aggregation, the Aggregator that the set of Aggregation Ports selects is the Aggregator with the same Port Number as the lowest-numbered Aggregation Port in the aggregation. Note that this lowest numbered Aggregation Port may not be in a state that allows data transfer across the link; however, it has selected the Aggregator in question. This is illustrated in Figure 6-20.




Figure 6-20—Selection of Aggregators

If the Aggregation Port is Individual, then the Aggregator selected is always the Aggregation Port's own Aggregator. Otherwise, an Aggregator is selected from the set of Aggregators corresponding to the set of Aggregation Ports that will form the aggregation. The Aggregator selected is the lowest numbered

1 Aggregator with the same selection parameters as those of the Aggregation Port. These selection parameters
2 are

- 3
- 4 a) The Actor's System ID
- 5 b) The Actor's operational Key
- 6 c) The Partner's System ID
- 7 d) The Partner's operational Key
- 8 e) The Individual_Aggregator state (which has to be FALSE)
- 9

10 **6.4.15 Mux machine**

11
12 The Mux machine shall implement the function specified in either of the Mux machine state diagrams,
13 Figure 6-21 and Figure 6-22, with their associated parameters (6.4.4 through 6.4.11).

14
15 The state machine conventions in IEEE Std 802.1Q-2011 require that all in-state actions are performed in
16 sequence and are atomic, completing before the evaluation of any exit conditions and before the execution
17 of procedures or evaluation of exit conditions for any other state or state machine.

18
19 The independent control state diagram (Figure 6-21) is suitable for use implementations in which it is
20 possible to control enabling and disabling of frame collection from an Aggregation Port, and frame
21 distribution to an Aggregation Port, independently. The coupled control state diagram (Figure 6-22) is
22 suitable for use implementations where collecting and distributing cannot be controlled independently with
23 respect to an Aggregation Port. It is recommended that the independent control state diagram be
24 implemented in preference to the coupled control state diagram.

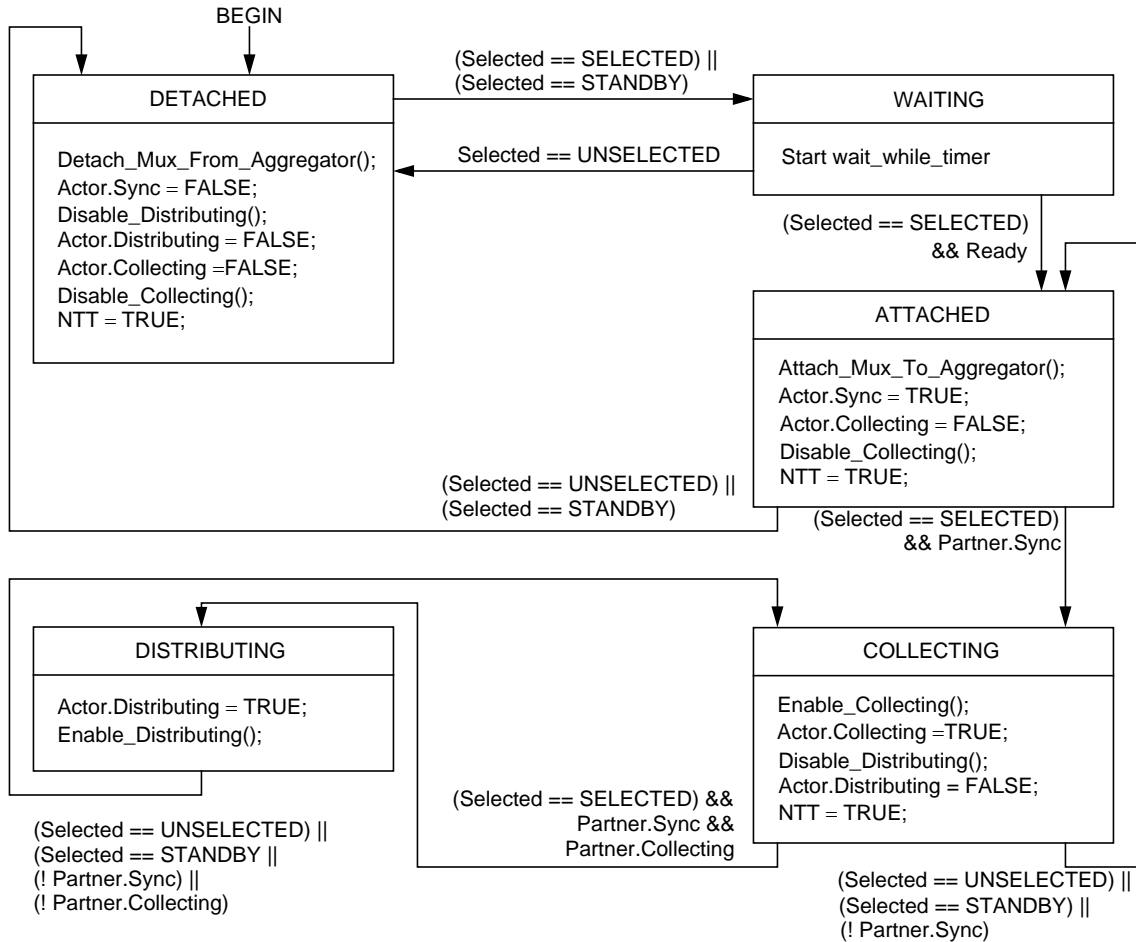
25
26 The value of the Selected variable may be changed by the following:

- 27
- 28 a) *The Receive machine.* The Receive machine can set Selected to UNSELECTED at any time if any of
29 the following change: The Partner System ID, the Partner Key, the Partner_State.Aggregation, the
30 Actor System ID, the Actor Key, or the Actor_State.Aggregation.
- 31 b) *The Selection Logic, in the process of selecting an Aggregator.* The Selection Logic will select an
32 Aggregator when the Mux machine is in the DETACHED state and the value of the Selected
33 variable is UNSELECTED.
- 34 c) *The Selection Logic, in the process of selecting or de-selecting standby links.* If the value of the
35 Selected variable is SELECTED or STANDBY, the Selection Logic can change the value to
36 STANDBY or SELECTED.
- 37

38 The Mux machine enters the WAITING state from the DETACHED state if the Selection Logic determines
39 that Selected is either SELECTED or STANDBY. The WAITING state provides a holding state for the
40 following two purposes:

- 41
- 42 d) If Selected is SELECTED, the wait_while_timer forces a delay to allow for the possibility that other
43 Aggregation Ports may be reconfiguring at the same time. Once the wait_while_timer expires, and
44 once the wait_while_timers of all other Aggregation Ports that are ready to attach to the same
45 Aggregator have expired, the process of attaching the Aggregation Port to the Aggregator can
46 proceed, and the state machine enters the ATTACHED state. During the waiting time, changes in
47 selection parameters can occur that will result in a re-evaluation of Selected. If Selected becomes
48 UNSELECTED, then the state machine re-enters the DETACHED state. If Selected becomes
49 STANDBY, the operation is as described in item e).

50 NOTE—This waiting period reduces the disturbance that will be visible to higher layers; for example, on start-
51 up events. However, the selection need not wait for the entire waiting period in cases where it is known that no
52 other Aggregation Ports will attach; for example, where all other Aggregation Ports with the same operational
53 Key are already attached to the Aggregator.



The following abbreviations are used in this diagram:
Actor.Sync: Actor_Oper_Port_State.Synchronization
Actor.Collecting: Actor_Oper_Port_State.Collecting
Actor.Distributing: Actor_Oper_Port_State.Distributing
Partner.Sync: Partner_Oper_Port_State.Synchronization
Partner.Collecting: Partner_Oper_Port_State.Collecting

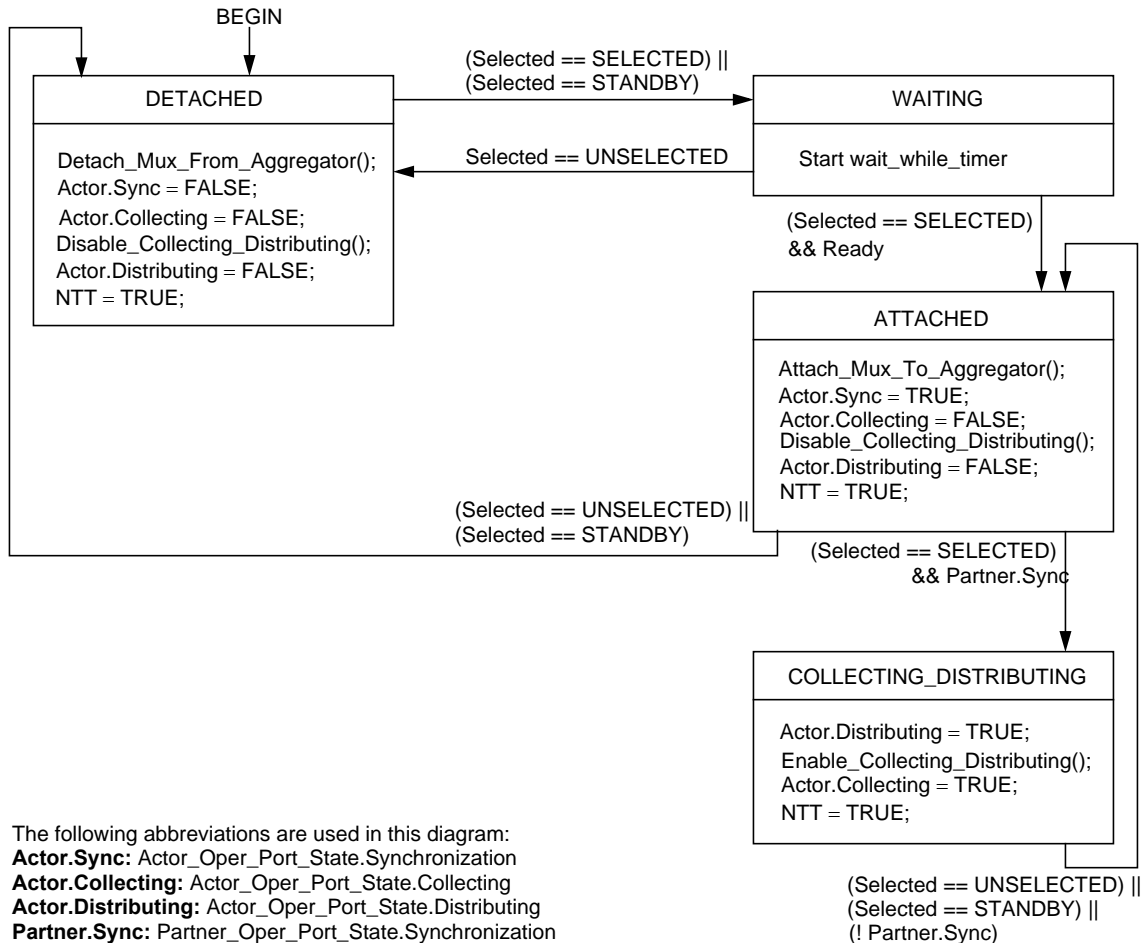
Figure 6-21—Mux machine state diagram (independent control)

- e) If Selected is STANDBY, the Aggregation Port is held in the WAITING state until such a time as the selection parameters change, resulting in a re-evaluation of the Selected variable. If Selected becomes UNSELECTED, the state machine re-enters the DETACHED state. If SELECTED becomes SELECTED, then the operation is as described in item d). The latter case allows an Aggregation Port to be brought into operation from STANDBY with minimum delay once Selected becomes SELECTED.

On entry to the ATTACHED state, the Mux machine initiates the process of attaching the Aggregation Port to the selected Aggregator. Once the attachment process has completed, the value of Actor_Oper_Port_State.Synchronization is set to TRUE indicating that the Actor considers the Aggregation Port to be IN_SYNC, and Actor_Oper_Port_State.Collecting is set to FALSE. Collection of frames from the Aggregation Port is disabled. In the coupled control state machine, Distribution of frames to the Aggregation Port is also disabled, and Actor_Oper_Port_State.Distributing is set to FALSE.

A change in the Selected variable to UNSELECTED or to STANDBY causes the state machine to enter the DETACHED state. The process of detaching the Aggregation Port from the Aggregator is started. Once the

1 detachment process is completed, Actor_Oper_Port_State.Synchronization is set to FALSE indicating that
 2 the Actor considers the Aggregation Port to be OUT_OF_SYNC, distribution of frames to the Aggregation
 3 Port is disabled, Actor_Oper_Port_State.Distributing and Actor_Oper_Port_State.Collecting are set to
 4 FALSE, and collection of frames from the Aggregation Port is disabled. The state machine remains in the
 5 DETACHED state until such time as the Selection logic is able to select an appropriate Aggregator.



37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Figure 6-22—Mux machine state diagram (coupled control)

While in the ATTACHED state, a TRUE value for Partner_Oper_Port_State.Synchronization causes the state machine to transition to the COLLECTING state (independent control) or the COLLECTING_DISTRIBUTING state (coupled control).

In the COLLECTING state, collection of frames from the Aggregation Port is enabled, then Actor_Oper_Port_State.Collecting is set to TRUE, then distribution of frames to the Aggregation Port is disabled and Actor_Oper_Port_State.Distributing is set to FALSE. The state machine will return to the ATTACHED state if Selected changes to UNSELECTED or STANDBY, or if the Partner's synchronization state becomes FALSE. Once the Partner has signaled that collecting has been enabled (Partner_Oper_Port_State.Collecting is TRUE), the state machine transitions to the DISTRIBUTING state. The value of Actor_Oper_Port_State.Distributing is set to TRUE, and then distribution of frames to the Aggregation Port is enabled. From DISTRIBUTING, a return to the COLLECTING state occurs if the value of Selected becomes UNSELECTED or STANDBY, if the Partner's synchronization state becomes FALSE, or if the Partner signals that collection has been disabled (Partner_Oper_Port_State.Collecting is FALSE).

1 In the COLLECTING_DISTRIBUTING state, the value of Actor_Oper_Port_State.Distributing is set to
2 TRUE, distribution of frames to the Aggregation Port and collection of frames from the Aggregation Port
3 are both enabled, and then Actor_Oper_Port_State.Collecting is set to TRUE. The state machine will return
4 to the ATTACHED state if Selected changes to UNSELECTED or STANDBY, or if the Partner's
5 synchronization state becomes FALSE.

6
7 The sequence of operations and transitions defined for the COLLECTING and DISTRIBUTING states in
8 the independent control version of this state machine ensures that frames are not distributed to an
9 Aggregation Port until the Partner has enabled collection, and that distribution is stopped as soon as the
10 Partner's state indicates that collection has been disabled. This sequence minimizes the possibility that
11 frames will be misdelivered during the process of bringing the Aggregation Port into operation or taking the
12 Aggregation Port out of operation. In the coupled control version of the state machine, the COLLECTING
13 and DISTRIBUTING states merge together to form the combined state, COLLECTING_DISTRIBUTING.
14 As independent control is not possible, the coupled control state machine does not wait for the Partner to
15 signal that collection has started before enabling both collection and distribution.

16
17 The NTT variable is set to TRUE in the DETACHED, ATTACHED, COLLECTING, and
18 COLLECTING_DISTRIBUTING states in order to ensure that the Partner is made aware of the changes in
19 the Actor's state variables that are caused by the operations performed in those states.

20 21 **6.4.16 Transmit machine**

22
23 When the Transmit machine creates a LACPDU for transmission, it shall fill in the following fields with the
24 corresponding operational values for this Aggregation Port:

- 25
26 a) Actor_Port and Actor_Port_Priority
27 b) Actor_System and Actor_System_Priority
28 c) Actor_Key
29 d) Actor_State
30 e) Partner_Port and Partner_Port_Priority
31 f) Partner_System and Partner_System_Priority
32 g) Partner_Key
33 h) Partner_State
34 i) CollectorMaxDelay

35
36 When the Periodic machine is in the NO_PERIODIC state, the Transmit machine shall

- 37
38 — Not transmit any LACPDUs, and
39 — Set the value of NTT to FALSE.

40
41 When the LACP_Enabled variable is TRUE and the NTT (6.4.7) variable is TRUE, the Transmit machine
42 shall ensure that a properly formatted LACPDU (6.4.2) is transmitted [i.e., issue a
43 CtrlMuxN:M_UNITDATA.Request(LACPDU) service primitive], subject to the restriction that no more
44 than three LACPDUs may be transmitted in any Fast_Periodic_Time interval. If NTT is set to TRUE when
45 this limit is in force, the transmission shall be delayed until such a time as the restriction is no longer in
46 force. The NTT variable shall be set to FALSE when the Transmit machine has transmitted a LACPDU.

47
48 If the transmission of a LACPDU is delayed due to the above restriction, the information sent in the
49 LACPDU corresponds to the operational values for the Aggregation Port at the time of transmission, not at
50 the time when NTT was first set to TRUE. In other words, the LACPDU transmission model is based upon
51 the transmission of state information that is current at the time an opportunity to transmit occurs, as opposed
52 to queuing messages for transmission.

1 When the LACP_Enabled variable is FALSE, the Transmit machine shall not transmit any LACPDU and
2 shall set the value of NTT to FALSE.
3

4 In a Version 2 implementation when enable_long_pdu_xmit and LongLACPDUtransmit are TRUE the
5 transmitted LACPDU will be a Long LACPDU, formatted as defined in 6.4.2 and including the Port
6 Conversation Mask TLVs (6.4.2.4.3).
7

8 **6.4.17 Churn Detection machines** 9

10 If implemented, the Churn Detection machines shall implement the functions specified in Figure 6-23 and
11 Figure 6-24 with their associated parameters (6.4.4 through 6.4.11). Implementation of the Churn Detection
12 machines is mandatory if the associated management functionality (the Aggregation Port Debug
13 Information package) is implemented; otherwise, implementation of the Churn Detection machines is
14 optional.
15

16 The Churn Detection machines detect the situation where an Aggregation Port is operable, but the Actor and
17 Partner have not attached the link to an Aggregator and brought the link into operation within a bounded
18 time period. Under normal operation of the LACP, agreement between Actor and Partner should be reached
19 very rapidly. Continued failure to reach agreement can be symptomatic of device failure, of the presence of
20 non-standard devices, or of misconfiguration; it can also be the result of normal operation in cases where
21 either or both Systems are restricted in their ability to aggregate. Detection of this condition is signaled by
22 the Churn Detection machines to management in order to prompt administrative action to further diagnose
23 and correct the fault.
24

25 NOTE—One of the classes of problems that will be detected by this machine is the one where the implementation has
26 been designed to support a limited number of Aggregators (fewer than the number of Aggregation Ports—see 6.7.4.2)
27 and the physical topology is such that one or more Aggregation Ports end up with no Aggregator to attach to. This may
28 be the result of a wiring error or an error in the allocation of operational Key values to the Aggregation Ports and
29 Aggregators. Alternatively, failure of a link to aggregate may be the result of a link being placed in standby mode by a
30 System that has hardware limitations placed on its aggregation ability, leading it to make use of the techniques described
31 in 6.7 to find the ideal configuration. The churn detection timers allow 60 seconds to elapse before a Churn condition is
32 signaled.
33

34 The symptoms that the Actor Churn Detection state machine detects is that the Actor's Mux has determined
35 that it is OUT_OF_SYNC, and that condition has not resolved itself within a period of time equal to
36 Short_Timeout_Time (6.4.4). Under normal conditions, this is ample time for convergence to take place.
37 Similarly, the Partner Churn Detection state machine detects a failure of the Partner's Mux to synchronize.
38

39 The Actor Churn Detection state machine is depicted in Figure 6-23.
40

41 The Partner Churn Detection state machine is depicted in Figure 6-24
42

43 **6.4.18 Long LACPDU machine** 44

45 The Long LACPDU machine shall implement the function specified in Figure 6-25 with its associated
46 parameters (6.4.4 through 6.4.11).
47

48 The Long LACPDU machine is only applicable to Version 2 or higher version implementations of LACP. It
49 enables the transmission on an Aggregation Port of Long LACPDU by an Actor implementing Version 2 or
50 higher LACP only if the LACPDU received on the same port report a Version 2 or higher implementation
51 of the LACP by the Partner.
52

53 The state machine has two states. They are as follows:
54

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

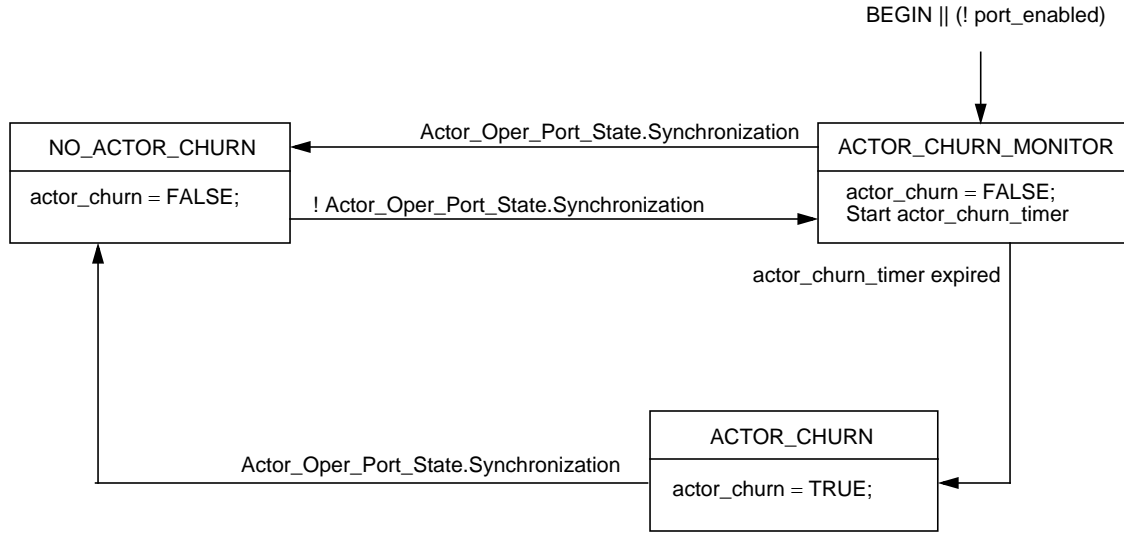


Figure 6-23—Actor Churn Detection machine state diagram

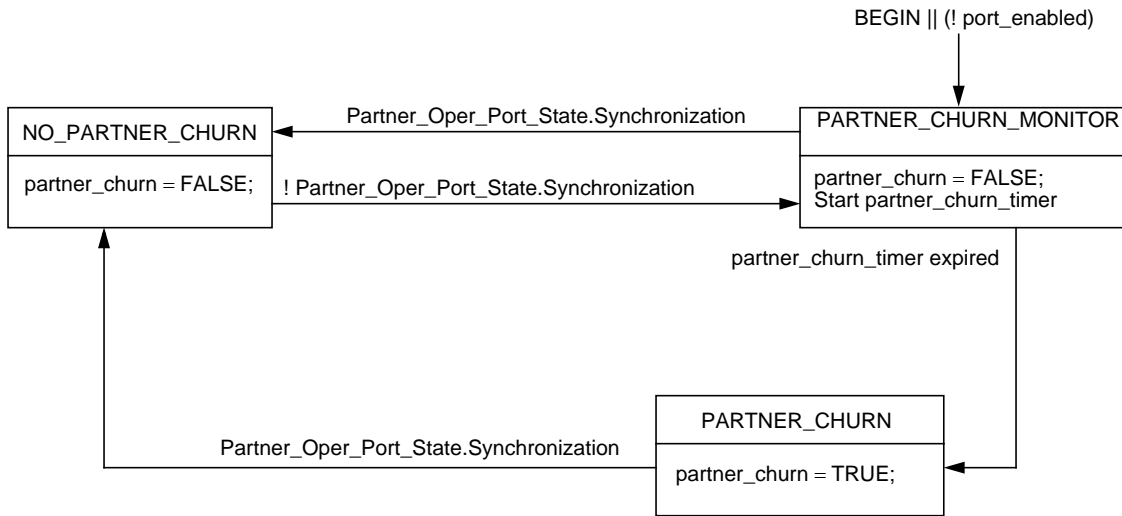
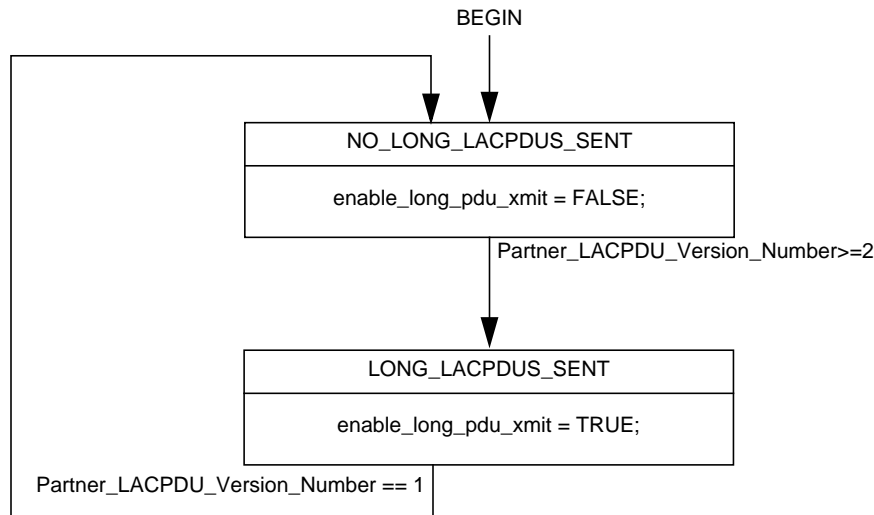


Figure 6-24—Partner Churn Detection machine state diagram

- a) *NO_LONG_LACPDUS_SENT*. While in this state, Long LACPDU transmission is not allowed.
- b) *LONG_LACPDUS_SENT*. While in this state, Long LACPDU transmission is allowed.

The System is initialized in the *NO_LONG_LACPDUS_SENT* and Long LACPDU transmission is prohibited. Upon reception of an LACPDU the Receive machine (6.4.12) transits to the *CURRENT* state and a Version 2 or higher implementation will run the recordVersionNumber function to update the Partner_LACPDU_Version_Number variable with the value in the Version Number field on the received LACPDU. If the updated value is larger or equal to 2 the System will transit to the *LONG_LACPDUS_SENT* state and transmission of Long LACPDU will be enabled. The System remains in the *LONG_LACPDUS_SENT* state until a received LACPDU reports a Version 1 implementation by the Partner System, which triggers a state transition to the *NO_LONG_LACPDUS_SENT* state and the transmission of Long LACPDU will be prohibited.



19
20
21
22
23
24
25

Figure 6-25—Long LACPDU machine state diagram

26 27 28 29 30 31 32 33 34 35 36 37 38 39 40

6.5 Marker protocol

41 42 43 44 45 46 47 48 49

6.5.1 Introduction

50
51
52
53
54

The Marker protocol allows the Frame Distribution function of an Actor's Link Aggregation sublayer to request the transmission of a Marker PDU on a given link. The Marker PDU is received by the Partner's Frame Collection function and a Marker Response PDU is returned on the same link to the initiating Actor's Frame Distribution function. Marker and Marker Response PDUs are treated by the underlying MACs at each end of the link as normal MPDUs; i.e., there is no prioritization or special treatment of Marker or Marker Response PDUs relative to other frames. Marker/Marker Response PDUs are subject to the operation of flow control, where supported on the link. Hence, if the Frame Distribution function requests transmission of a Marker PDU on a given link and does not transmit any further MPDUs that relate to a given set of conversations until the corresponding Marker Response PDU is received from that link, then it can be certain that there are no MPDUs related to those conversations still to be received by the Partner's Frame Collection function. The use of the Marker protocol can therefore allow the Frame Distribution function a means of determining the point at which a given set of conversations can safely be reallocated from one link to another without the danger of causing frames in those conversations to be misordered at the Frame Collector.

NOTE—The use of the Marker protocol is further discussed in Annex B. An alternative to the Marker protocol is defined in 6.6.

The operation of the Marker protocol is unaffected by any changes in the Collecting and Distributing states associated with the Aggregation Port. Therefore, Marker and Marker Response PDUs can be sent on an Aggregation Port whose Frame Distribution function is disabled; similarly, such PDUs can be received and passed to the relevant Aggregator's Frame Collection or Distribution function on an Aggregation Port whose Frame Collection function is disabled.

The use of the Marker protocol is optional; however, the ability to respond to Marker PDUs, as defined for the operation of the Marker Responder (see 6.5.4.1 and 6.5.4.2), is mandatory. Some distribution algorithms may not require the use of a marker; other mechanisms (such as timeouts) may be used as an alternative.

This standard does not specify how, when, or if the Marker protocol is used. It is possible to define a distribution algorithm that does not require the Marker Protocol, e.g., that defined in 8.2. See also Annex B.

The Marker protocol does not provide a guarantee of a response from the Partner; no provision is made for the consequences of frame loss or for the failure of the Partner System to respond correctly. Implementations that make use of this protocol have to therefore make their own provision for handling such cases.

6.5.2 Sequence of operations

Figure 6-26 illustrates the sequence of marker operations between an initiating and responding System. Time is assumed to flow from the top of the diagram to the bottom.




Figure 6-26—Marker protocol time sequence diagram

6.5.3 Marker and Marker Response PDU structure and encoding

6.5.3.1 Transmission and representation of octets

All Marker and Marker Response PDUs comprise an integral number of octets. The bits in each octet are numbered from 0 to 7, where 0 is the low-order bit. When consecutive octets are used to represent a numerical value, the most significant octet is transmitted first, followed by successively less significant octets.

When the encoding of (an element of) a Marker or Marker Response PDU is depicted in a diagram, then

- Octets are transmitted from top to bottom.
- Within an octet, bits are shown with bit 0 to the left and bit 7 to the right, and are transmitted from left to right.
- Numerical values are encoded as binary numbers.

6.5.3.2 Encapsulation of Marker and Marker Response PDU in frames

Marker and Marker Response PDUs are encoded in the `mac_service_data_unit` parameter of an `M_UNITDATA.request` or `M_UNITDATA.indication`. The first octets of the `mac_service_data_unit` are a protocol identifier, followed by the LACPDU, followed by padding octets, if any, as required by the underlying MAC service.

Where the ISS instance used to transmit and receive frames is provided by a media access control method that can support EtherType encoding directly (e.g., is an IEEE 802.3 MAC), the protocol identifier is two octets in length, and the value is the Slow Protocols EtherType (hexadecimal 88-09).

Where the ISS instance is provided by a media access method that cannot directly support EtherType encoding (e.g., is an IEEE 802.11 MAC), the TPID is encoded according to the rule for a Subnetwork Access Protocol (Clause 10 of IEEE Std 802) that encapsulates Ethernet frames over LLC, and comprises the SNAP header (hexadecimal AA-AA-03) followed by the SNAP PID (hexadecimal 00-00-00) followed by the Slow Protocols EtherType (hexadecimal 88-09).

6.5.3.3 Marker and Marker Response PDU structure

The Marker PDU and Marker Response PDU structure shall be as shown in Figure 6-27 and as further described in the following field definitions:

- a) *Subtype*. The Subtype field identifies the specific Slow Protocol being encapsulated. Both Marker and Marker Response PDUs carry the Marker_subtype value 0x02.
- b) *Version number*. This identifies the Marker protocol version; implementations conformant to this version of the standard carry the value 0x01.
- c) *TLV_type = Marker Information/Marker Response Information*. This indicates the nature of the information carried in this TLV-tuple. Marker Information is encoded as the integer value 0x01; Marker Response Information is encoded as the integer value 0x02.
- d) *Marker_Information_Length/Marker_Response_Information_Length*. This field indicates the length (in octets) of this TLV-tuple. Both Marker and Marker Response information use a length value of 16 (0x10).
- e) *Requester_Port*. The Port Number assigned to the Aggregation Port by the Requester (the System sending the initial Marker PDU), encoded as an unsigned integer.

Marker PDU	Octets	Marker Response PDU
Subtype = Marker	1	Subtype = Marker
Version Number	1	Version Number
TLV_type = Marker Information	1	TLV_type = Marker Response Information
Marker_Information_Length= 16	1	Marker_Response_Information_Length = 16
Requester_Port	2	Requester_Port
Requester_System	6	Requester_System
Requester_Transaction_ID	4	Requester_Transaction_ID
Pad = 0	2	Pad = 0
TLV_type = Terminator	1	TLV_type = Terminator
Terminator_Length = 0	1	Terminator_Length = 0
Reserved	90	Reserved
FCS	4	FCS

Figure 6-27—Marker PDU and Marker Response PDU structure

- f) *Requester_System*. The MAC address component of the Requester's System ID.
- g) *Requester_Transaction_ID*. The transaction ID allocated to this request by the requester, encoded as an integer.
- h) *Pad*. This field is used to align TLV-tuples on 16-byte memory boundaries. It is transmitted as zeros in Marker PDUs; in Marker Response PDUs, this field may be transmitted as zeros, or with the contents of this field from the Marker PDU triggering the response. The field is ignored on receipt in all cases.

NOTE 1—The difference in handling of the Pad field in Marker Response PDUs allows an implementation to reflect the contents of the received Marker PDU in its response, without enforcing the requirement to transmit the field as zeros.

- i) *TLV_type = Terminator*. This field indicates the nature of the information carried in this TLV-tuple. Terminator (end of message) information carries the integer value 0x00.
- j) *Terminator_Length*. This field indicates the length (in octets) of this TLV-tuple. Terminator information uses a length value of 0 (0x00).
- k) *Reserved*. These 90 octets are reserved for use in future extensions to the protocol. It is transmitted as zeros in Marker PDUs; in Marker Response PDUs, this field may be either transmitted as zeros, or with the contents of this field from the Marker PDU triggering the response. The field is ignored on receipt in all cases.

NOTE 2—These trailing reserved octets are included in all Marker and Marker Response PDUs in order to force a fixed PDU size, regardless of the version of the protocol. Hence, a Version 1 implementation is guaranteed to be able to receive version N PDUs successfully, although version N PDUs may contain additional information that cannot be interpreted (and will be ignored) by the Version 1 implementation. A crucial factor in ensuring backwards compatibility is that any future version of the protocol is required not to redefine the structure or semantics of information defined for the previous version; it may only add new information elements to the previous set. Hence, in a version N PDU, a Version 1 implementation can expect to find the Version 1 information in exactly the same places as in a Version 1 PDU, and can expect to interpret that information as defined for Version 1.

- l) *FCS*. This field is the Frame Check Sequence, typically generated by the underlying MAC.

6.5.4 Protocol definition

6.5.4.1 Operation of the marker protocol

Marker PDUs may be generated by the Frame Distribution function to provide a sequence marker in the stream of frames constituting a conversation or set of conversations. Received Marker PDUs are delivered to the Marker Responder within the Frame Collection function of the Partner System.

On receipt of a valid Marker PDU, the Frame Collection function issues a Marker Response PDU, in the format specified in Figure 6-27, to the same Aggregation Port from which the Marker PDU was received. The **Requester_Port**, **Requester_System**, and **Requester_Transaction_ID** parameter in the Marker Response PDU carry the same values as those received in the corresponding Marker PDU.

Received Marker Response PDUs are passed to the Marker Receiver within the Frame Distribution function. Implementation of the Marker Generator and Receiver is optional.

The Marker Generator, if implemented, shall comply with the frame rate limitation constraint for Slow Protocols, as specified in IEEE Std 802.3 Annex 57A.3. A Marker Responder may (but is not required to) control its Marker Response transmissions to conform to this Slow Protocols timing constraint when faced with Marker messages not in compliance with this constraint (i.e., to send fewer Marker Response PDUs than Marker PDUs received). If the Marker Responder is controlling its responses in this manner, Marker Response PDUs corresponding to Marker PDUs received in excess of the Slow Protocols timing constraint shall not be sent.

NOTE 1—It is important that Marker Response PDUs not be queued indefinitely, and sent long after the corresponding Marker PDU that triggered the response.

Frames generated by the Marker Responder do not count towards the rate limitation constraint for Slow Protocols, as specified in IEEE Std 802.3 Annex 57A.3.

NOTE 2—The Marker Responder behaves the same whether Link Aggregation is employed in a single system or as part of a Distributed Resilient Network Interconnect (DRNI, Clause 9). In the latter case, frames in transit between one Portal System and another when a Marker Response is generated can be delivered out of order. See 6.6 for an alternative method of ensuring frame ordering.

6.5.4.2 Marker Responder state diagram

The Marker Responder shall implement the function specified in Figure 6-28, with its associated parameters (6.5.4.2.1 through 6.5.4.2.2).

6.5.4.2.1 Variables

DA

SA

mac_service_data_unit

priority

The parameters of the M_UNITDATA.request and M_UNITDATA.indication primitives.

Protocol_DA

One of the addresses selected from Table 6-1 determined by the setting of the aAggPortProtocolDA managed object (7.3.2.2.1).

Value: 48 bits.

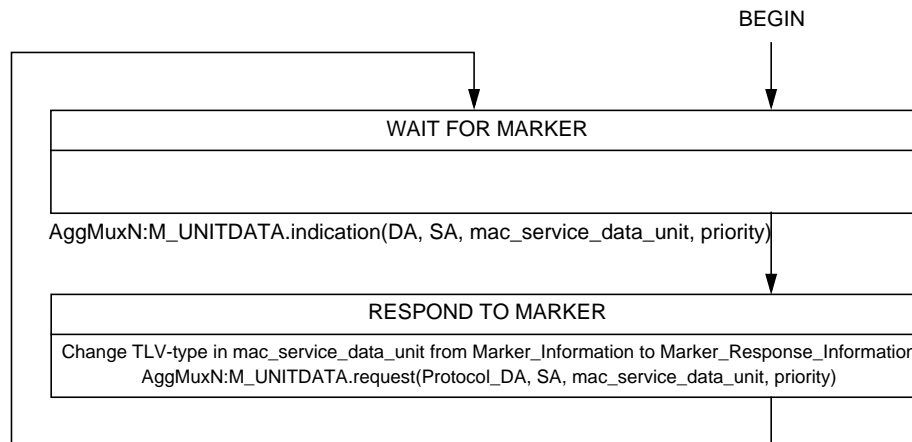
6.5.4.2.2 Messages

AggMuxN:M_UNITDATA.request

The service primitive used to transmit a frame with the specified parameters.

AggMuxN:M_UNITDATA.indication

The service primitive used to pass a received frame to a client with the specified parameters.



The value of N (the Port Number) in the AggMuxN:M_UNITDATA.request primitive shall be the same as that of the received AggMuxN:M_UNITDATA.indication

Figure 6-28—Marker Responder state diagram

Upon receipt of an AggMuxN:M_UNITDATA.indication primitive, the Marker Responder shall not validate the Version Number, Pad, or Reserved fields in the contained Marker Request PDU. The same actions are taken regardless of the values received in these fields. A Marker Responder may validate the Marker_Information_Length field. These behaviors, together with the constraint on future protocol

1 enhancements discussed in the Note in 6.5.3.3 item n), allow Version 1 devices to be compatible with future
2 revisions of the protocol.

6.6 Conversation-sensitive frame collection and distribution

3
4
5
6
7 Conversation-sensitive frame collection and distribution provides a means for both the Frame Distributor
8 and the Frame Collector to determine which Aggregation Link is to be selected by the Distributor for any
9 given frame, and for the Collector to only accept frames received on the expected Aggregation Link. This
10 serves a dual function. It enables the use of the same Aggregation Link for both directions of a bidirectional
11 conversation, and it provides a means of ensuring that frames are not misordered when a conversation is
12 moved between Aggregation Links without requiring the use of the Marker Protocol (6.5).

13
14 The use of the Marker protocol has the advantage that a single, simple Frame Collector can be used to
15 receive frames from any Frame Distributor, no matter what distribution algorithm is used. However,
16 ensuring that a Marker Response PDU works properly in a DRNI (Clause 9), where either or both ends of
17 the Link Aggregation Group can comprise multiple Systems, becomes problematical.

18
19 A Frame Distributor and Frame Collector at the two ends of a Link Aggregation Group can implement
20 Conversation-sensitive frame collection and distribution. This includes the following elements:

- 21
22 a) TLVs are included in LACPDUs to indicate and acknowledge the distribution algorithm in use by
23 the Frame Distributor.
- 24 b) The Frame Distributor assigns every frame to a Port Conversation ID (8.1), and uses that Port
25 Conversation ID to assign the frame to an Aggregation Port.
- 26 c) The Frame Collector knows and can execute the same frame-to-Conversation ID mapping that is
27 used by the Frame Distributor.
- 28 d) TLVs included in LACPDUs allow the Frame Distributor to specify which Conversation IDs are
29 being transmitted on each Aggregation Port. The Frame Collector stores this information in the
30 `Collection_Conversation_Mask`.
- 31 e) The Frame Collector discards frames received on Aggregation Ports whose Conversation IDs are
32 not in the list supplied by the Frame Distributor for that Aggregation Port.

33
34 Thus, frame ordering is guaranteed by the fact that the Frame Collector receives a given conversation on
35 only one Aggregation Port. In the case of a DRNI (Clause 9), the Frame Collectors in a Portal are configured
36 to ensure that a given conversation is received from only one Aggregation Port.

37
38 NOTE—Whether Conversation-sensitive frame collection and distribution is required of a given conformant System
39 depends on the support of per service frame distribution algorithms by that System. (See Clause 8).

40
41 The Frame Collector extracts a frame's Port Conversation ID according to the distribution algorithm
42 employed (e.g., that specified in 8.2), then examines the `Collection_Conversation_Mask` variable (6.6.1.1.2)
43 of the Aggregation Port on which the frame was received. The frame is passed on to the Client Port only if
44 `Collection_Conversation_Mask` passes its Port Conversation ID.

6.6.1 Conversation-sensitive collection and distribution state diagrams

45
46
47 When a System implements Conversation-sensitive collection and distribution:

- 48
49 a) The Frame Distributor shall operate as specified in 6.2.4 and implement the state diagram shown in
50 Figure 6-4 and the associated definitions contained in 6.2.4.1, constrained though to use only the
51 frame distribution, configured in the `aAggPortAlgorithm` (7.3.1.1.33) and the
52 `aAggConversationAdminLink[]` (7.3.1.1.35), and provided by `Port_Oper_Conversation_Mask`
53 (6.6.2.2) in the case that the Conversation-sensitive LACP operation (6.6.2) is supported.

- b) The Frame Collector shall implement the function specified in the Conversation-sensitive collection state diagram shown in Figure 6-29 and the associated definitions contained in 6.6.1.1.

6.6.1.1 Conversion-sensitive collection state diagram

6.6.1.1.1 Variables

DA

SA

mac_service_data_unit

priority

The parameters of the M_UNITDATA.indication primitive.

BEGIN

A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-)initialization has completed.

Value: Boolean

Discard_Wrong_Conversation

Discard frames with incorrect Port Conversation IDs flag. If TRUE, the Frame Collector discards any frame received from an Aggregation Port with a Port Conversation ID not set in Collection_Conversation_Mask and the Frame Distributor uses only the frame distribution algorithm selected by the aAggPortAlgorithm (7.3.1.1.33). If FALSE, the Frame Collector does not filter frames received from the Aggregation Ports. The variable is set equal to aAggAdminDiscardWrongConversation (7.3.1.1.37).

Value: Boolean

6.6.1.1.2 Variables associated with each Aggregation Port

Collection_Conversation_Mask

The operational Boolean vector, indexed by Port Conversation ID, indicating whether the indexed Port Conversation ID is allowed to reach the Aggregator when received through this Aggregation Port (TRUE = passes). This variable is constructed from the agreed Conversation_PortList[], and if supported, the operation of the Conversation-sensitive LACP (6.6.2).

Value: sequence of Boolean values, indexed by Port Conversation ID.

6.6.1.1.3 Functions

DeterminePortConversationID

This function determines a Port Conversation ID for the frame. Any algorithm can be used to determine the Port Conversation ID provided it depends only on information found within the frame and the same algorithm is used by both the Frame Distributor and Frame Collector. This function may involve the application of local Service ID to Conversation ID mappings, provided by 7.3.1.1.38, in cases where additional local information is needed to determine the Port Conversation ID, as is the case of the Per I-SID service distribution (8.2.3).

6.6.1.1.4 Messages

Agg:M_UNITDATA.indication

AggMuxN:M_UNITDATA.indication

The service primitives used to pass a received frame to a client with the specified parameters.

6.6.1.1.5 State diagram

The Conversation-sensitive collection state diagram shall implement the function specified in Figure 6-29 with its associated parameters (6.6.1.1.1 through 6.6.1.1.4)

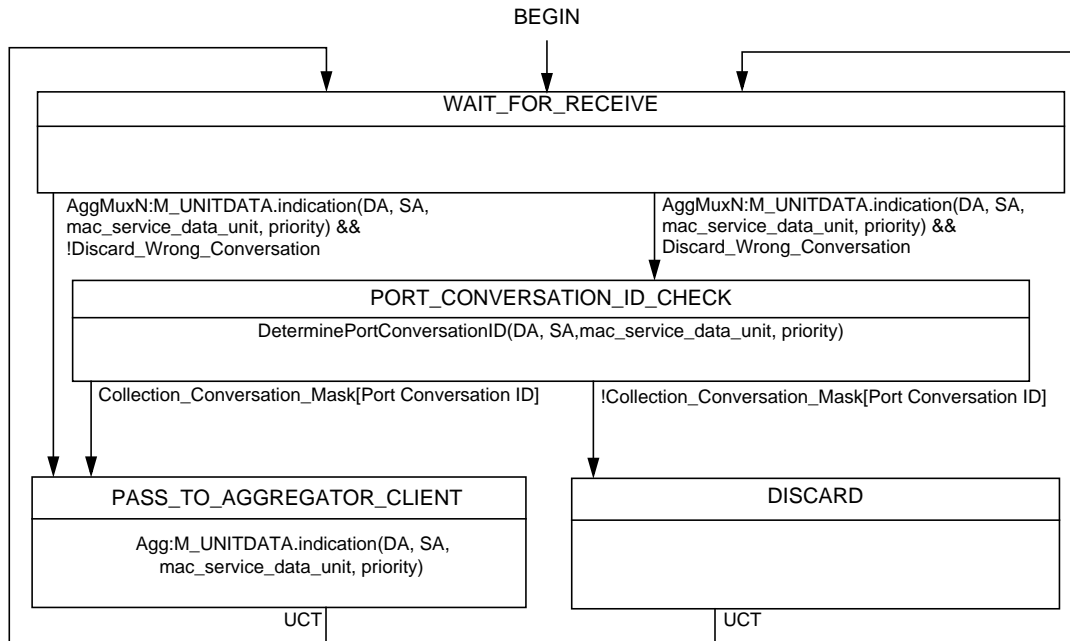


Figure 6-29—Conversation-sensitive collection state diagram

The operation of the Frame Collector when Conversation-sensitive collection and distribution is implemented is identical to that specified in 6.2.3 when the variable `Discard_Wrong_Conversation` is `FALSE`. When `Discard_Wrong_Conversation` is `TRUE`, indicating that the Conversation-sensitive collection and distribution function is operational, the Frame Collector operation is modified so that the Frame Collector goes into the `PORT_CONVERSATION_ID_CHECK` state and the `DeterminePortConversationID` function is run to determine a Port Conversation ID for the frame. This value is used as an index to the `Collection_Conversation_Mask` variable and if the associated identified Boolean is `TRUE` the frame passes to the Aggregator Client, otherwise it gets discarded.

6.6.2 Conversation-sensitive LACP state diagrams

The Conversation-sensitive LACP Version 2 additions to the Version 1 LACP specified in 6.4 enable configuration, control and coordination of the Frame Collection and the Frame Distributor functions of each participating System or Portal by using static information local to each System or Portal and dynamic information exchanged by means of the Version 2 LACPDUs exchanges specified in this clause.

For each Aggregation Port in a System or Portal, Conversation-sensitive LACP:

- a) Maintains the agreed configuration information;
- b) Exchanges configuration information with the Partner System to verify agreed per Port Conversation ID configuration choices;
- c) Uses information from the Partner System's Link Aggregation Control entity to enable or disable the Aggregator's Frame Collector and Frame Distributor on a Port Conversation ID basis.

1 The operation of Conversation-sensitive LACP is described in detail in 6.6.2.1 through 6.6.2.7. The
2 Conversation-sensitive LACP protocol entity is optional protocol sub-entity within the Link Aggregation
3 Control in Figure 6-2.
4

5 **6.6.2.1 Per-Aggregator Variables**

6 Actor_Conversation_LinkList_Digest

7 A digest of aAggConversationAdminLink[] (7.3.1.1.35) for exchange with the Partner System.
8 The digest, is a 16-octet MD5 fingerprint [see IETF RFC 1321 (1992)] created from the
9 aAggConversationAdminLink[]. To calculate the digest, aAggConversationAdminLink[] is
10 considered to contain 4096 consecutive elements, where each element contains a list of Link
11 Number IDs, encoded as binary numbers in the order of preference, highest to lowest, followed
12 by the Port Conversation ID. The first element of the table contains the prioritized list of Link
13 Number IDs of Aggregation Ports assigned to Port Conversation ID 0, the second element the
14 prioritized list of Link Number IDs of Aggregation Ports assigned to Port Conversation ID 1,
15 the third element the prioritized list of Link Number IDs of Aggregation Ports assigned to Port
16 Conversation ID 2, and so on, with the last element containing the prioritized list of Link
17 Number IDs of Aggregation Ports assigned to Port Conversation ID 4095. This variable is
18 referenced by the Port Conversation ID Digest TLV (6.4.2.4.2).
19 Value: MD5 digest

20 Actor_Conversation_Service_Mapping_Digest

21 A digest of aAggAdminServiceConversationMap[] for exchange with the Partner System. The
22 digest, is a 16-octet MD5 fingerprint [see IETF RFC 1321 (1992)] created from the
23 aAggAdminServiceConversationMap[]. To calculate the digest,
24 aAggAdminServiceConversationMap[] is considered to contain 4096 consecutive elements,
25 where each element contains, in general a set of Service IDs (8.2.2) encoded as binary numbers
26 in increasing order followed by the Port Conversation ID to which they are mapped. If the
27 Service IDs are representing VIDs, only a single 12-bit VID is mapped to each service, while in
28 the case that Service IDs are representing I-SIDs, more than one 24-bit I-SIDs are possible. The
29 first element of the table contains the Service ID assigned to Port Conversation ID 0, the
30 second element the Service ID assigned to Port Conversation ID 1, the third element the
31 Service ID assigned to Port Conversation ID 2, and so on, with the last element containing the
32 Service ID assigned to Port Conversation ID 4095. This variable is referenced by the optional
33 Port Conversation Service Mapping TLV (6.4.2.4.4) when the Service IDs are different from
34 the Port Conversation IDs and a table is associating their values.
35 Value: MD5 digest

36 Actor_Port_Algorithm

37 The algorithm used by the Actor to assign frames to Port Conversation IDs. Always set equal to
38 aAggPortAlgorithm (7.3.1.1.33). This variable is referenced by the Port Algorithm TLV
39 (6.4.2.4.1)

40 Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining
41 this algorithm followed by one octet identifying this specific algorithm).
42

43 Conversation_PortList[]

44 This array holds the mapping of Port Conversation IDs to Aggregation Ports inside the System
45 or Portal, derived from the primary aAggConversationAdminLink[] managed object. An array
46 of 4096 lists, indexed by Conversation ID, that determines which Aggregation Port in this
47 System (or Portal) carries which Conversation ID. Each item in the array is a list of the Port IDs
48 (6.3.4) of the Aggregation Ports in this System (or in this Portal, if this Aggregator is attached
49 to a DR Function [9.3]), in priority order from most desired to least desired, for carrying the
50 indexed Conversation ID. This variable is updated by the updateConvesationPortList function
51 which is always invoked when a new aAggConversationAdminLink[] (7.3.1.1.35) or new
52 aAggPortLinkNumberID (7.3.2.1.27) operator command is issued.

53 Value: sequence of Port IDs (6.3.4)

54 Differ_Conversation_Service_Digests

1 A Boolean indicating that the Conversation Service Mapping Digests used by the Systems or
2 Portals at the two ends of the LAG are different.
3 Value: Boolean
4 Differ_Port_Algorithms
5 A Boolean indicating that the Port Algorithms used by the Systems or Portals at the two ends of
6 the LAG are different.
7 Value: Boolean
8 Differ_Port_Conversation_Digests
9 A Boolean indicating that the Port Conversation Digests used by the Systems or Portals at the
10 two ends of the LAG are different.
11 Value: Boolean
12 Partner_Admin_Conversation_PortList_Digest
13 The value for the digest of the prioritized Port Conversation ID-to-Link Number ID
14 assignments of the Partner System, assigned by administrator or System policy for use when
15 the Partner's information is unknown. Its default value is set to NULL.
16 Value: MD5 digest
17 Partner_Admin_Conversation_Service_Mapping_Digest
18 The value for the digest of the Port Conversation ID-to-Service ID assignments of the Partner
19 System, assigned by administrator or System policy for use when the Partner's information is
20 unknown. Its default value is set to NULL.
21 Value: MD5 digest
22 Partner_Admin_Port_Algorithm
23 The value for the algorithm of the Partner System, assigned by administrator or System policy
24 for use when the Partner's information is unknown. Its default value is set to NULL.
25 Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining
26 this algorithm followed by one octet identifying this specific algorithm).
27 Partner_Conversation_PortList_Digest
28 The last-received digest of the Partner System's prioritized Port Conversation ID-to-
29 Aggregation Port assignments.
30 Value: MD5 digest
31 Partner_Conversation_Service_Mapping_Digest
32 The last-received digest of the Partner System's Port Conversation ID-to-Service ID
33 assignments.
34 Value: MD5 digest
35 Partner_Port_Algorithm
36 The last-received value of the algorithm used by the Partner System to assign frames to Port
37 Conversation IDs.
38 Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining
39 this algorithm followed by one octet identifying this specific algorithm).
40

41 **6.6.2.2 Variables associated with each Aggregation Port**

42
43 Admin_Link_Number_ID
44 The Link Number ID value configured for this Aggregation Port by the System's administrator.
45 It is always set equal to aAggPortLinkNumberID (7.3.2.1.27).
46 Value: Integer, 0 to 65535
47 ChangeActorOperDist
48 This variable tracks the variable Actor_Oper_Port_State.Distributing. It sets the WTR_timer to
49 aAggPortWTRTime when the Actor_Oper_Port_State.Distributing changes from FALSE to
50 TRUE and there is at least one Aggregation Port in the same LAG having a non NULL
51 Port_Oper_Conversation_Mask. The variable is set to TRUE when the WTR_timer is running
52 and the Actor_Oper_Port_State.Distributing changes from TRUE to FALSE or when the
53 WTR_timer expires and Actor_Oper_Port_State.Distributing is TRUE. This variable is also set
54 to TRUE if new values for the aAggConversationAdminLink[] (7.3.1.1.35) or

1 aAggAdminServiceConversationMap[] (7.3.1.1.38) are initiated. When the variable is set to
2 TRUE, updateLocal is set to TRUE on all other Aggregation Ports attached to the same
3 Aggregator. ChangeActorOperDist is set to FALSE by the updateConversationMask function.
4 Value: Boolean

5 ActPar_Sync
6 This variable indicates if the Conversation Mask used by the Actor's Frame Distributor is the
7 same or not as that used by the Partner Frame Distributor in the same System. TRUE if
8 Partner_Oper_Conversation_Mask == Port_Oper_Conversation_Mask, otherwise FALSE.
9 This variable is referenced by the Port Conversation Mask-1 TLV (6.4.2.4.3).
10 Value: Boolean

11 Comp_Oper_Conversation_Mask
12 The operational value of the Port_Oper_Conversation_Mask Boolean vector of the System or
13 Portal as computed by updateConversationMask.
14 Value: sequence of Boolean values, indexed by Port Conversation ID.

15 Link_Number_ID
16 This System's operational Link Number ID value for this Aggregation Port. It is used to assign
17 the Aggregation Port to the Link Number ID which is an identifier commonly used by this
18 System's administrator and the Partner System's administrator to refer to the interconnected
19 link.
20 Value: Integer, 0 to 65535

21 LongLACPDUtransmit
22 This variable indicates to the Transmit machine (6.4.16) that it needs to transmit Long
23 LACPDU's including the Port Conversation Mask TLVs (6.4.2.4.3). It is set to TRUE while the
24 current_while_Long_LACP_timer runs and set to FALSE otherwise.
25 Value: Boolean

26 NTT
27 Need To Transmit flag. This variable is re-used from the basic LACP operation (6.4.7).
28 Value: Boolean
29 TRUE indicates that there is new protocol information that should be transmitted on the
30 link, or that the Partner needs to be reminded of the old information.
31 FALSE otherwise.

32 Partner_Admin_Conversation_Mask
33 The Port_Oper_Conversation_Mask Boolean vector of the Partner System, assigned by
34 administrator or System policy for use when the Partner's information is unknown. Its default
35 is set to NULL.
36 Value: sequence of Boolean values, indexed by Port Conversation ID.

37 Partner_Admin_Link_Number_ID
38 The value for the Link Number ID of the Partner System for this Aggregation Port, assigned by
39 administrator or System policy for use when the Partner's information is unknown. It is always
40 equal to aAggPortPartnerAdminLinkNumberID (7.3.2.1.28).
41 Value: Integer, 0 to 65535.

42 Partner_ActPar_Sync
43 The last-received value of the Partner System's ActPar_Sync value.
44 Value: Boolean

45 Partner_DWC
46 The last-received value of the Partner System's DWC value.
47 Value: Boolean

48 Partner_Link_Number_ID
49 The last-received value of the Partner System's Link Number ID value.
50 Value: Integer, 0 to 65535

51 Partner_Oper_Conversation_Mask
52 The last received Boolean vector indicating whether the indexed Port Conversation ID can be
53 distributed through the Partner Systems' Aggregation Port on the link (TRUE = passes).
54 Value: sequence of Boolean values, indexed by Port Conversation ID.

1 Partner_PSI
 2 The last-received value of the Partner System's PSI value.
 3 Value: Boolean
 4 Port_Oper_Conversation_Mask
 5 The operational Boolean vector, indexed by Port Conversation ID, indicating whether the
 6 indexed Port Conversation ID is distributed through this Aggregation Port (TRUE = passes).
 7 This variable is constructed from the agreed Conversation_PortList[], and the operation of the
 8 Conversation-sensitive LACP (6.6.2). This variable is split in four parts,
 9 Port_Oper_Conversation_Mask_1, Port_Oper_Conversation_Mask_2,
 10 Port_Oper_Conversation_Mask_3, and Port_Oper_Conversation_Mask_4 in order to be
 11 carried by the Port Conversation Mask TLVs as specified in 6.4.2.4.3.
 12 Value: sequence of Boolean values, indexed by Port Conversation ID.
 13 updateLocal
 14 This variable is set to TRUE by the recordReceivedConversationMaskTLV function when the
 15 value of either ActPar_Sync or Partner_ActPar_Sync is FALSE.
 16

17 6.6.2.3 Variables used for managing the operation of the state diagrams

18
 19 BEGIN
 20 This variable indicates the initialization (or reinitialization) of the Conversation-sensitive
 21 LACP protocol entity. It is set to TRUE when the System is initialized or reinitialized, and is
 22 set to FALSE when (re-)initialization has completed.
 23 Value: Boolean
 24 actor_CDS_churn
 25 This variable indicates that the Actor CDS Churn Detection machine has detected that the
 26 Conversation Masks used by Frame Distributor and the Frame Collector in the local System
 27 have failed to converge within a specified time, and that management intervention is required.
 28 Value: Boolean
 29 ChangeAggregationPorts
 30 This variable tracks the operational state of all Aggregation Ports associated to this System and
 31 is set to TRUE when any of them changes. It is the logical OR of the ChangeActorOperDist
 32 variables for all Aggregation Ports.
 33 Value: Boolean
 34 partner_CDS_churn
 35 This variable indicates that the Partner CDS Churn Detection machine has detected that the
 36 Conversation Masks used by Frame Distributor and the Frame Collector in the remote System
 37 have failed to converge within a specified time, and that management intervention is required.
 38 Value: Boolean
 39 updateAllLocal
 40 This variable is the logical OR of the updateLocal variables for all Aggregation Ports.
 41 Value: Boolean
 42

43 6.6.2.4 Functions

44
 45 recordReceivedConversationMaskTLV
 46 This function records the parameter value for the *ActPar_Sync* carried in a received Port
 47 Conversation Mask-1 TLV (6.4.2.4.3) as the current operational parameter values for the
 48 Partner_ActPar_Sync, it concatenates the value of *Port_Oper_Conversation_Mask_1*,
 49 *Port_Oper_Conversation_Mask_2*, *Port_Oper_Conversation_Mask_3*, and
 50 *Port_Oper_Conversation_Mask_4* carried by the Port Conversation Mask-1 TLV, ..., etc.
 51 (6.4.2.4.3), and it records the concatenation as the current value for the
 52 Partner_Oper_Conversation_Mask variable. If this function is implemented by a Portal System
 53 (Clause 9), it also records the parameter value for the *PSI* carried in a received Port
 54

1 Conversation Mask-1 TLV (6.4.2.4.3) as the current operational parameter value for the
2 Partner_PSI.
3
4 It then compares the variable Port_Oper_Conversation_Mask to the
5 Partner_Oper_Conversation_Mask and, if they are different:
6 It sets ActPar_Sync to FALSE;
7 It sets updateLocal to TRUE;
8 Else:
9 It sets ActPar_Sync to TRUE;
10 If Partner_ActPar_Sync is FALSE, sets updateLocal to TRUE.
11 recordConversationPortDigestTLV
12 This function records the parameter value for the *Link_Number_ID* and the
13 *Actor_Conversation_LinkList_Digest* carried in a received Port Conversation ID Digest TLV
14 (6.4.2.4.2) as the current operational parameter value for the Partner_Link_Number_ID and the
15 Partner_Conversation_PortList_Digest respectively.
16
17 It compares the newly updated value of the Partner_Link_Number_ID to the value of
18 Link_Number_ID and if they are different, then:
19 The reportToManagement function is invoked to alert the Manager to the existence of a
20 configuration error in the Aggregation ports selection choices for the Port Conversation
21 IDs between the Systems or Portals at the ends of the LAG. In addition if this System's or
22 Portal's Identifier is numerically larger than the Partner's, then
23 Link_Number_ID == Partner_Link_Number_ID; and
24 The function updateConvesationPortList is invoked.
25
26 It then compares the newly updated value of the Partner_Conversation_PortList_Digest to the
27 value of Actor_Conversation_LinkList_Digest and, if they are different:
28 The variable Differ_Port_Conversation_Digests is set to TRUE.
29 Otherwise:
30 The variable Differ_Port_Conversation_Digests is set to FALSE.
31 recordConversationServiceMappingDigestTLV
32 This function records the parameter value for the
33 *Actor_Conversation_Service_Mapping_Digest* carried in a received Port Conversation Service
34 Mapping TLV (6.4.2.4.4) as the current operational parameter value for
35 Partner_Conversation_Service_Mapping_Digest.
36
37 It then compares the newly updated value of the
38 Partner_Conversation_Service_Mapping_Digest to the value of
39 Actor_Conversation_Service_Mapping_Digest and, if they are different:
40 The variable Differ_Conversation_Service_Digests is set to TRUE.
41 Otherwise:
42 The variable Differ_Conversation_Service_Digests is set to FALSE.
43 recordDefaultConversationMask
44 This function sets Partner_Oper_Conversation_Mask and Collection_Conversation_Mask to
45 Partner_Admin_Conversation_Mask and sets ActPar_Sync to FALSE.
46 recordDefaultConversationPortDigest
47 This function records the default parameter values for the Partner as follows:
48 Partner_Conversation_PortList_Digest
49 == Partner_Admin_Conversation_PortList_Digest;
50 Partner_Link_Number_ID == Partner_Admin_Link_Number_ID.
51 It also initializes Link_Number_ID to its configured value as follows:
52 Link_Number_ID == Admin_Link_Number_ID
53 recordDefaultConversationServiceMappingDigest
54

1 This function records the default parameter value for the Partner provided by the
2 Partner_Admin_Conversation_Service_Mapping_Digest as the current Partner operational
3 parameter value for Partner_Conversation_Service_Mapping_Digest.

4 recordDefaultPortAlgorithm

5 This function records the default parameter value for the Partner provided by the
6 Partner_Admin_Port_Algorithm as the current Partner operational parameter value for
7 Partner_Port_Algorithm.

8 recordPortAlgorithmTLV

9 This function records the parameter value for the *Actor_Port_Algorithm* carried in a received
10 Port Algorithm TLV (6.4.2.4.1) as the current operational parameter value for
11 Partner_Port_Algorithm.

12
13 It then compares the newly updated value of the Partner_Port_Algorithm to the value of
14 Actor_Port_Algorithm and if they are different or any of them has the value 0 (Table 6-4):

15 The variable Differ_Port_Algorithms is set to TRUE.

16 Otherwise:

17 The variable Differ_Port_Algorithms is set to FALSE.

18 reportToManagement

19 To alert the Manager to the existence of a configuration error in the algorithm or agreed
20 Aggregation ports selection choices for the Port Conversation IDs between the Systems or
21 Portals at the ends of the LAG.

22 resetAllChangeActorOperDist

23 To reset ChangeActorOperDist to FALSE on all ports.

24 resetAllupdateLocal

25 To reset updateLocal to FALSE on all ports.

26 updateConversationMask

27 This function computes a new value for the Port_Oper_Conversation_Mask based on the
28 Conversation_PortList[].

29
30 It first computes a new Boolean vector for the Comp_Oper_Conversation_Mask to be used by
31 this Aggregation Port as follows: For every indexed Port Conversation ID, the selection
32 priority list for all Aggregation Ports in the same LAG, which is provided by
33 Conversation_PortList[], is checked to identify and exclude all Aggregation Ports for which the
34 Actor_Oper_Port_State.Distributing = FALSE (condition that covers the cases where the
35 associated Aggregation Ports are either non operable (port_enabled = FALSE), or in an
36 EXPIRED state, or not in the LAG) and all Aggregation Ports for which the WTR_timer is
37 running. The new Boolean vector for the Comp_Oper_Conversation_Mask to be used by this
38 Aggregation Port is computed based on this modified priority list by setting TRUE for every
39 indexed Port Conversation ID that has the highest selection priority for this Aggregation Port.
40 If this function is implemented by a Portal System (Clause 9), with its
41 DRF_Portal_System_Number value set to a value that is different than 1, its Portal's System
42 Identifier set to a value that is numerically lower than the Partner's System Identifier, and
43 PSI == Partner_PSI == TRUE, then Comp_Oper_Conversation_Mask is set to NULL. Such
44 functionality is set to prohibit network partition when the Portal Systems in both interconnected
45 Portals become isolated.

46
47 It then compares the value of Port_Oper_Conversation_Mask to the new value of
48 Comp_Oper_Conversation_Mask. If they are different:

49 It updates Collection_Conversation_Mask as follows:

50 If ChangeAggregationPorts == TRUE or, updateAllLocal == TRUE, or, when the
51 function is implemented by a Portal System, IppAllUpdate == TRUE;

52 It sets the Collection_Conversation_Mask to the Boolean vector corresponding
53 to the result of the logical AND operation between the
54

1 Comp_Oper_Conversation_Mask and the current
2 Collection_Conversation_Mask;
3 Otherwise if ChangeAggregationPorts == FALSE and updateAllLocal == TRUE
4 and, when the function is implemented by a Portal System, IppAllUpdate == FALSE;
5 It sets Collection_Conversation_Mask = Comp_Oper_Conversation_Mask;
6 It updates the Port_Oper_Conversation_Mask to the new value of
7 Comp_Oper_Conversation_Mask;
8 Otherwise:
9 it leaves Port_Oper_Conversation_Mask and Collection_Conversation_Mask unmodified.

11 Finally it sets ChangeActorOperDist to FALSE and if ActPar_Sync or Partner_ActPar_Sync is
12 FALSE it starts the current_while_Long_LACP_timer, using Slow_Periodic_Time (6.4.4) as
13 the start value; sets NTT to TRUE; and keeps the variable LongLACPDUtransmit TRUE until
14 the timer's expiration.

15 updateConvesationPortList

16 This function transforms the mapping of Port Conversation IDs to Link Number IDs
17 configured by the aAggConversationAdminLink[] managed object into the equivalent mapping
18 of Port Conversation IDs to Port IDs used inside the System or Portal. The function updates
19 Conversation_PortList[] based on the current aAggConversationAdminLink[] and the
20 operational Link_Number_ID values. The resulted Conversation_PortList[] provides the
21 operational Aggregation Port agreed selection prioritized list per Port Conversation ID
22 expressed in terms of Port IDs. The function is always invoked when a new command to
23 change aAggConversationAdminLink[] or aAggPortLinkNumberID is issued.

25 6.6.2.5 Timers

27 current_while_Long_LACP_timer

28 This timer is used to indicate the transmission of Long LACPDU. It is always started with the
29 value Slow_Periodic_Time (6.4.4).

30 actor_CDS_churn_timer

31 This timer is used to detect Actor CDS churn states. It is started using the value
32 Churn_Detection_Time (see 6.4.4).

33 partner_CDS_churn_timer

34 This timer is used to detect Partner CDS churn states. It is started using the value
35 Churn_Detection_Time (see 6.4.4).

36 WTR_timer

37 This timer is used to prevent frequent distribution changes due to an intermittent defect. It
38 allows for a fixed period of time to elapse when the Actor_Oper_Port_State.Distributing
39 changes from FALSE to TRUE before updating the Port_Oper_Conversation_Mask. It is
40 started using the value aAggPortWTRTime (7.3.2.1.29). There is one per Aggregation Port.

42 The timers specified in this subclause have an implementation tolerance of ± 250 ms.

44 6.6.2.6 Messages

46 CtrlMuxN:M_UNITDATA.indication(LLACPDU)

47 This message is generated by the Control Parser as a result of the reception of a Long
48 LACPDU, formatted as defined in 6.4.2.

50 6.6.2.7 State diagrams

52 The Conversation-sensitive LACP state diagrams shall implement the functions specified in Figure 6-31,
53 Figure 6-32, Figure 6-33, and Figure 6-34, and may implement the functions in Figure 6-35 and

Figure 6-36, with their associated parameters (6.6.2.1 through 6.6.2.6 and in 6.4). There are six Conversation-sensitive LACP state diagrams for every Aggregation Port in the LAG:

- a) The Verification state diagram (*VER*) depicted in Figure 6-31;
- b) The Report for Management Action state diagram (*RMA*) depicted in Figure 6-32;
- c) The Receive Long LACPDU state diagram (*RXL*) depicted in Figure 6-33;
- d) The Update Mask state diagram (*UM*) depicted in Figure 6-34;
- e) The Actor Collection and Distribution Sensitive (CDS) Churn Detection state machine is depicted in Figure 6-35;
- f) The Partner Collection and Distribution Sensitive (CDS) Churn Detection state machine is depicted in Figure 6-36.

Figure 6-30 illustrates the relationships among these state machines and the flow of information between them.

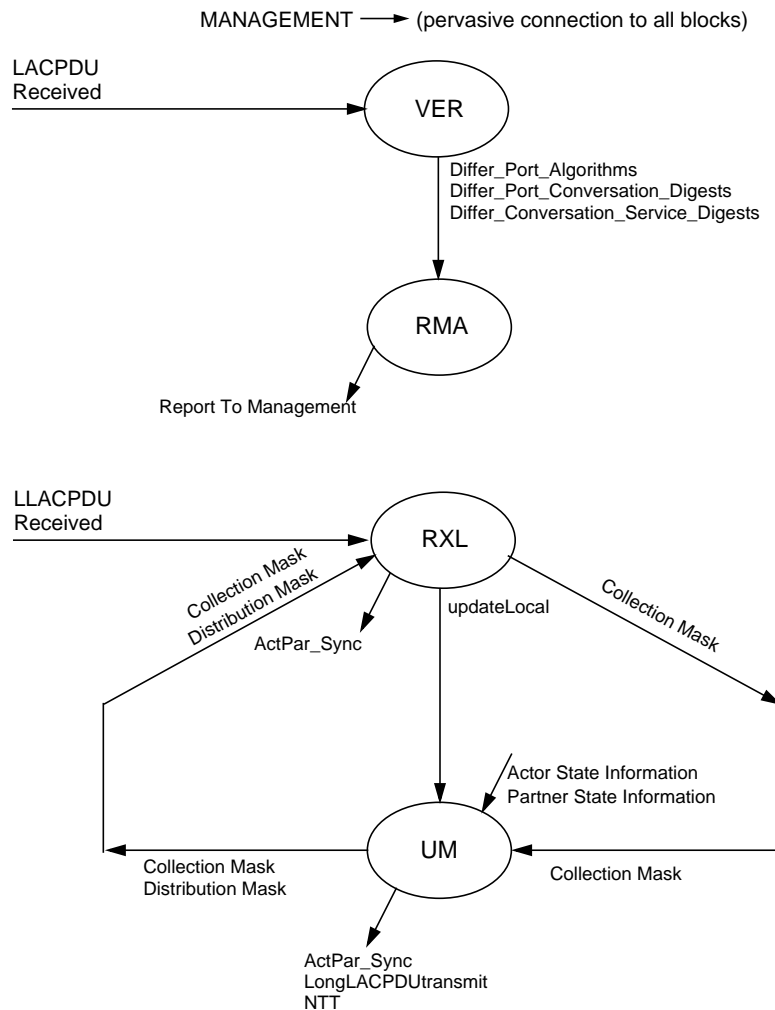


Figure 6-30—Interrelationships among Conversation-sensitive LACP state machines

The set of arrows labeled Distribution Mask represents the Port_Oper_Conversation_Mask, contained in transmitted Long LACPDU or supplied by administrative default values, and being updated by the Update Mask machine. The set of arrows labeled Collection Mask represents the flow of updated

Collection_Conversation_Mask values between the RXL and UM state machines. The set of arrows labeled Partner State Information and Actor State information are as defined in 6.4.3. Partner State Information represents new Partner information, contained in an incoming LACPDU or supplied by administrative default values, being fed by the RX machine. Actor State Information represents the flow of updated Actor state information as provided by the LACP state machines. Transmission of LACPDUs occurs either as a result of the Periodic machine determining the need to transmit a periodic LACPDU, or as a result of changes to the Actor's state information that need to be communicated to the Partner. The need to transmit a LACPDU is signaled to the Transmit machine by asserting NTT. The remaining arrows represent shared variables in the state machine description that allow a state machine to cause events to occur in another state machine.

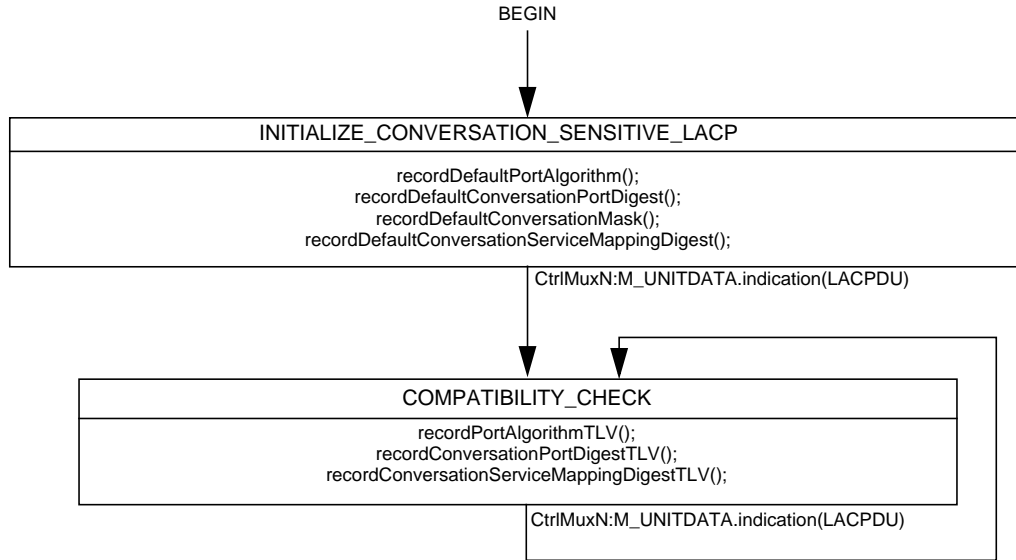


Figure 6-31—Verification state diagram

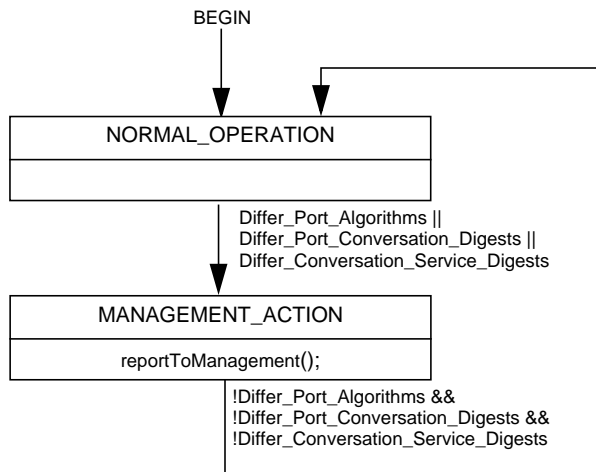


Figure 6-32—Report for Management Action state diagram

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

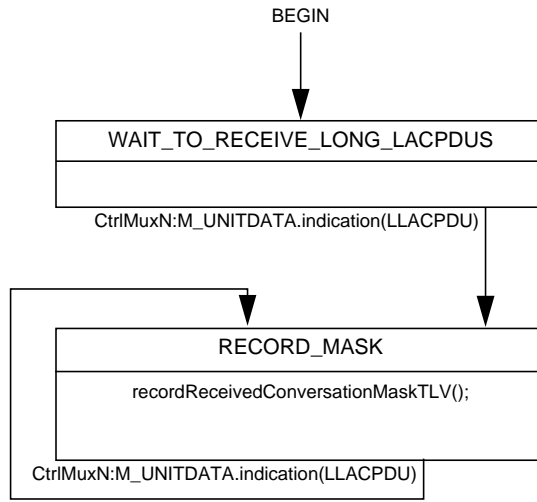


Figure 6-33—Receive Long LACPDU state diagram

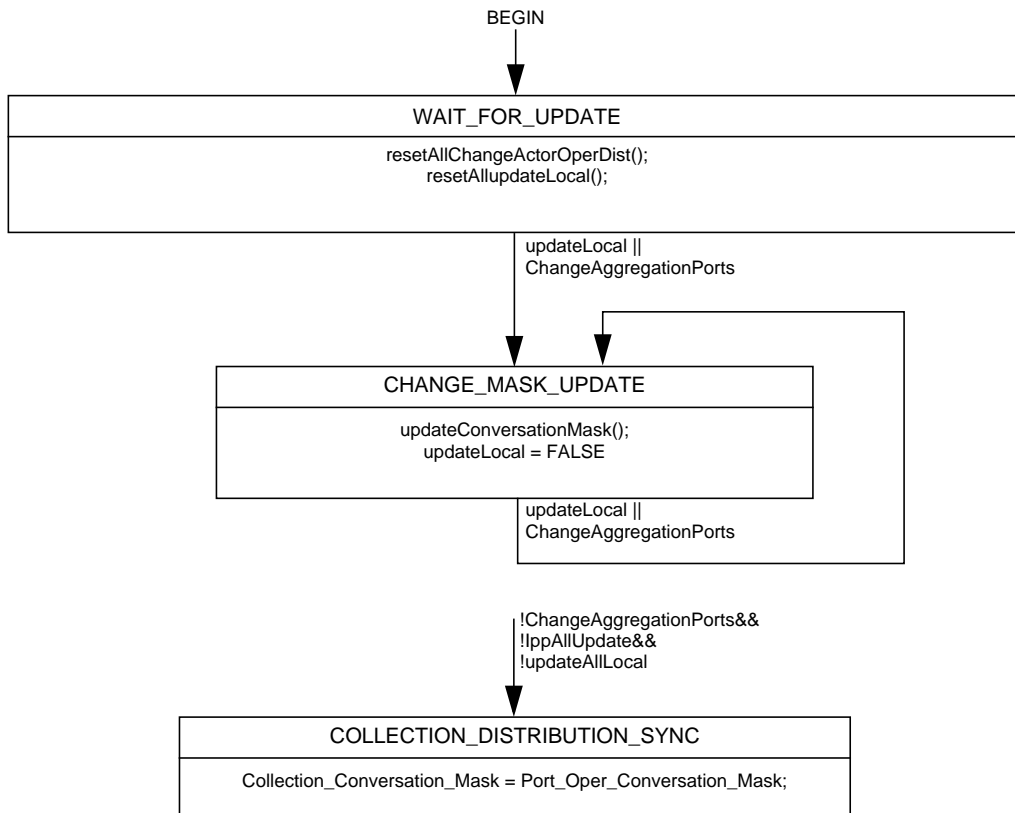


Figure 6-34—Update Mask state diagram

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

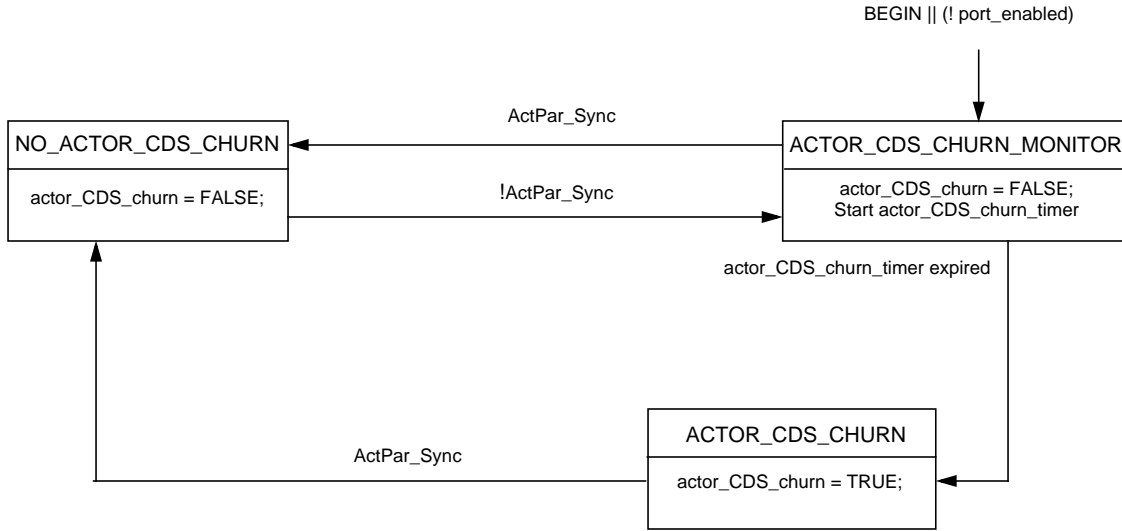


Figure 6-35— Actor CDS Churn Detection machine state diagram

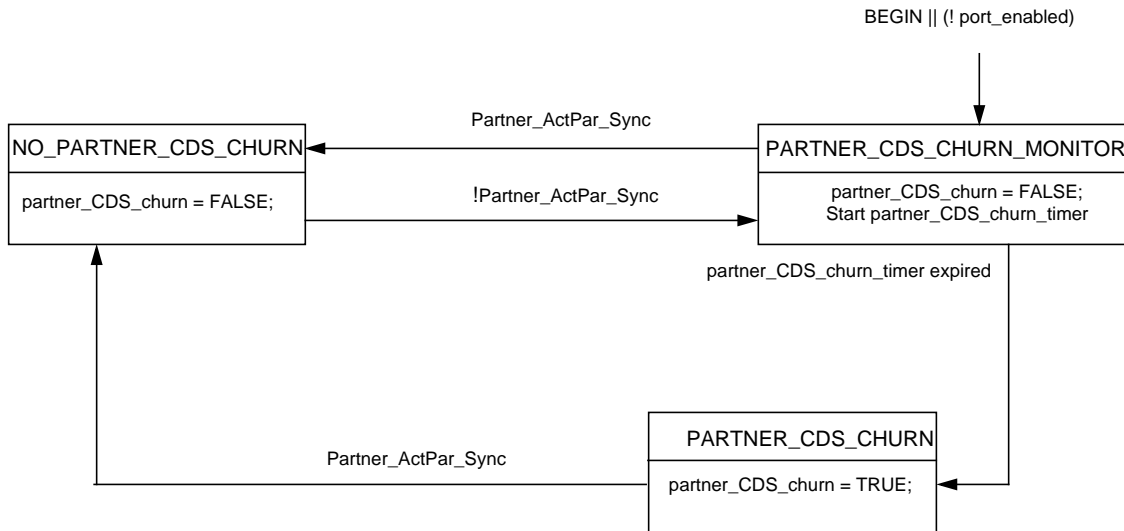


Figure 6-36—Partner CDS Churn Detection machine state diagram

These state diagrams are not required unless both Systems or Portals are running a Version 2 implementation of LACP. Through the basic LACP operation (6.4), a LAG is established between an agreed set of Aggregation Ports which are configured to use the same Operational Keys to their selected Aggregator. All Aggregation Ports in each participating System or Portal are using the same Operational Keys to their selected Aggregator. In addition, the two Systems or Portals have already agreed, by some administrative

1 means, to use a common algorithm to assign frames to Port Conversation IDs and this information has been
2 configured in both Systems or Portals (7.3.1.1.33). Furthermore, the two Systems or Portals have agreed on
3 the same selection priority list of Aggregation Ports, for each Port Conversation ID, and based on this, each
4 System can configure itself in order to accept and transmit the agreed services on their selected Aggregation
5 Ports as follows.

6
7 On Initialization, both Systems or Portals are configuring their Version 2 TLVs, with default values. The
8 Port Algorithm TLV (6.4.2.4.1) is configured with the agreed algorithm (7.3.1.1.33), an MD5 digest of
9 aAggConversationAdminLink[] (7.3.1.1.35) is used to configure the Port Conversation ID Digest TLV
10 (6.4.2.4.2), while in cases where frames are distributed to physical links tagged with Service IDs being
11 different than the Port Conversation IDs (8.2.3), an MD5 digest of aAggAdminServiceConversationMap[]
12 (7.3.1.1.38) is used to configure the Port Conversation Service Mapping TLV (6.4.2.4.4).

13
14 Reception of an LACPDU when in the INITIALIZE_CONVERSATION_SENSITIVE_LACP state triggers
15 compatibility tests. If the algorithm carried in the received Port Algorithm TLV (6.4.2.4.1), and/or the set of
16 reported Port Conversation IDs carried in the received Port Conversation ID Digest TLV (6.4.2.4.2), and/or
17 the Service ID to Port Conversation IDs mappings carried in the received Port Conversation Service
18 Mapping TLV (6.4.2.4.4), are different than the associated expected configurations, the state machine
19 transits to the MANAGEMENT_ACTION and the management system is notified in order to take further
20 actions such as to report on the Partner System's inconsistencies, to initialize the operation of the Marker
21 protocol, or other actions that are outside the scope of this standard.

22
23 Any event that changes one or more of the following:

- 24 — the operational status of any of the Aggregation Ports in the LAG;
- 25 — the configured selection rules provided by aAggConversationAdminLink[] (7.3.1.1.35);
- 26 — the Service ID to Port Conversation IDs provided by aAggAdminServiceConversationMap[]
27 (7.3.1.1.38);

28
29 will trigger a transition to the CHANGE_MASK_UPDATE state and (re-)calculation of the Conversation
30 Masks. The updateConversationMask() computes a Distribution Conversation Mask, which is provided by
31 the Port_Oper_Conversation_Mask variable, and a Collection Conversation Mask for the Aggregation port
32 based on the agreed selection rules, and starts the current_while_Long_LACP_timer using
33 Slow_Periodic_Time as the start value, and keeps the variable LongLACPDUtransmit TRUE until the
34 timer's expiration which causes the Transmit machine to transmit Long LACPDU whenever an NTT
35 trigger is set. The Collection Conversation Mask will always be set equal to the Distribution Conversation
36 Mask when the Distribution Conversation Masks of all Aggregation Ports in the same System have been
37 updated as a result of the local change. In the case of a Portal (Clause 9) such condition includes updates of
38 the IPP Conversation Masks.

39
40 On receipt of a Long LACPDU the receiving Aggregation Port stores the Partner's distribution Conversation
41 Mask carried in received Port Conversation Mask TLVs (6.4.2.4.3), and compares it with its own. If they
42 differ, or the port on the other end of the LAG is reporting that the Conversation Masks are not in sync,
43 further Long LACPDU exchanges will be triggered until both Aggregation Ports on the two ends of the
44 LAG start reporting the same set of Conversation Masks. If no differences are reported, the
45 current_while_Long_LACP_timer will terminate and the subsequent LACPDU exchanges will not include
46 the Port Conversation Mask TLVs.

47
48 The CDS Churn Detection machines detect the situation where the Actor and Partner Conversation Masks
49 cannot synchronize within a bounded time period. Under normal operation of the Conversation-Sensitive
50 LACP, agreement between Actor and Partner on the operational Conversation Masks used by the Frame
51 Distributor and the Frame Collector will be reached rapidly. Continued failure to reach agreement can be
52 symptomatic of device failure, of the presence of non-standard devices, or of misconfiguration. Detection of
53
54

1 this condition is signaled by the CDS Churn Detection machines to management in order to prompt
2 administrative action to further diagnose and correct the fault.

3
4 The Actor CDS Churn Detection state machine detects a failure of the Actor's and Partner's Conversation
5 Mask to converge. Under normal conditions, this is ample time for convergence to take place. Similarly, the
6 Partner Churn Detection state machine detects a failure on the Partner's view of the Conversation Masks to
7 converge.

8
9 NOTE—Administrative control of the port_enabled variable on a Aggregation Port through its associated
10 MAC_Enabled variable (Clause 11.2 of IEEE Std 802.1AC-2012) provides means for manual switch capabilities per
11 Aggregation Port in conversation-sensitive LACP. The operation of manual switch on a Port Conversation ID basis is
12 similar to that of changing the configured priority list (7.3.1.1.34).

13 14 15 **6.7 Configuration capabilities and restrictions**

16 17 **6.7.1 Use of system and port priorities**

18
19 The Link Aggregation Group identifier format specified by this standard cannot represent two or more Link
20 Aggregation Groups that share the same combination of {SK, TL}. The use of static keys in combination
21 with size restriction (e.g., for a Key Group containing six members, restricting the size of any aggregation to
22 four members or fewer) leads to one Link aggregation Group.

23
24 In Systems that have limited aggregation capability of this form, the following algorithm shall be used to
25 determine the subset of Aggregation Ports that will be aggregated together:

- 26
27 a) The System Aggregation Priority of each System is an eight octet binary number, formed by using
28 the Actor_System_Priority as the two most significant octets and the Actor's MAC address as the
29 least significant six octets. For a given Actor and Partner, the System with the numerically lower
30 value of System Aggregation Priority has the higher priority.
- 31 b) The Port Aggregation Priority of each Aggregation Port is a four octet binary number, formed by
32 using the Actor_Port_Priority as the two most significant octets and the Port Number as the two
33 least significant octets. For any given set of Aggregation Ports, the Aggregation Port with the
34 numerically lower value of Port Aggregation Priority has the higher priority.
- 35 c) Aggregation Ports shall be selected for aggregation by each System based upon the Port
36 Aggregation Priority assigned by the System with the higher System Aggregation Priority, starting
37 with the highest priority Aggregation Port of the System with the higher priority, and working
38 downward through the ordered list of Port Aggregation Priority values for the N Aggregation Ports,
39 applying the particular constraints imposed on the System concerned.
- 40 d) For each link that a given System cannot include in the aggregation, the Selection Logic identifies
41 the Selection state of the corresponding Aggregation Port as STANDBY, preventing the link from
42 becoming active. The synchronization state signaled in transmitted LACPDUs for such links will be
43 OUT_OF_SYNC.
- 44 e) The selection algorithm is reapplied upon changes in the membership of the Link Aggregation
45 Group (for example, if a link fails, or if a new link joins the group) and any consequent changes to
46 the set of active links are made accordingly.

47
48 The use of the standby capability is discussed in Annex C.

49 50 **6.7.2 Dynamic allocation of operational Keys**

51
52 In some circumstances, the use of System and port priorities may prove to be insufficient to generate the
53 optimum aggregation among the set of links connecting a pair of Systems. A System may have a limited
54 aggregation capability that cannot be simply expressed as a limit on the total number of links in the

1 aggregation. The full description of its restrictions may be that it can only aggregate together particular
2 subsets of links, and the sizes of the subsets need not all be the same.

3
4 NOTE 1—An example would be an implementation organized such that, for a set of four links A through D, it would be
5 possible to operate with {A+B+C+D} as a single aggregation, or operate with {A+B} and {C+D} as two separate
6 aggregations, or operate as four individual links; however, all other aggregation possibilities (such as {A+C} and
7 {B+D}) would not be achievable by the implementation.

8
9 In such circumstances, it is permissible for the System with the higher System Aggregation Priority (i.e., the
10 numerically lower value) to dynamically modify the operational Key value associated with one or more of
11 the Aggregation Ports; the System with the lower priority shall not attempt to modify operational Key values
12 for this purpose. Operational Key changes made by the higher priority System should be consistent with
13 maintaining its highest priority Aggregation Port in the aggregate as an active link (i.e., in the IN_SYNC
14 state). Successive operational Key changes, if they occur, should progressively reduce the number of
15 Aggregation Ports in the aggregation. The original operational Key value should be maintained for the
16 highest priority Aggregation Port thought to be aggregateable.

17
18 NOTE 2—Restricting operational Key changes in the manner described prevents the case where both Partner Systems
19 involved have limited capability and both attempt to make operational Key changes; this could be a non-converging
20 process, as a change by one participant can cause the other participant to make a change, which in turn causes the first
21 participant to make a change—and so on, ad infinitum.

22
23 This approach effectively gives the higher priority System permission to search the set of possible
24 configurations, in order to find the best combination of links given its own and its Partner's configuration
25 constraints. The reaction of the Partner System to these changes can be determined by observing the changes
26 in the synchronization state of each link. A System performing operational Key changes should allow at
27 least 4 s for the Partner System to change an OUT_OF_SYNC state to an IN_SYNC state.

28
29 In the course of normal operation an Aggregation Port can dynamically change its operating characteristics
30 (e.g., data rate, point-to-point operation). It is permissible (and appropriate) for the operational Key value
31 associated with such an Aggregation Port to change with the corresponding changes in the operating
32 characteristics of the link, so that the operational Key value always correctly reflects the aggregation
33 capability of the link. Operational Key changes that reflect such dynamic changes in the operating
34 characteristics of a link may be made by either System without restriction.

35 36 **6.7.3 Link Aggregation on shared-medium links**

37
38 The Link Aggregation Control Protocol cannot detect the presence of multiple Aggregation-aware devices
39 on the same link. Hence, shared-medium links shall be treated as Individual, with transmission/reception of
40 LACPDUs disabled on such Aggregation Ports.

41 42 **6.7.4 Selection Logic variants**

43
44 Two variants of the Selection Logic rules are described as follows:

- 45
46 a) The first accommodates implementations that may wish to operate in a manner that minimizes
47 disturbance of existing aggregates, at the expense of the deterministic characteristics of the logic
48 described in 6.4.14.2.
- 49 b) The second accommodates implementations that may wish to limit the number of Aggregators that
50 are available for use to fewer than the number of Aggregation Ports supported.

6.7.4.1 Reduced reconfiguration

By removing the constraint that the Aggregator chosen is always the lowest numbered Aggregator associated with the set of Aggregation Ports in an aggregation, an implementation can minimize the degree to which changes in the membership of a given aggregation result in changes of connectivity at higher layers. As there would still be the same number of Aggregators and Aggregation Ports with a given operational Key value, any Aggregation Port will still always be able to find an appropriate Aggregator to attach to, however the configuration achieved over time (i.e., after a series of link disconnections, reconnections, or reconfigurations) with this relaxed set of rules would not necessarily be the same as the configuration achieved if all Systems involved were reset, given the rules stated in 6.4.14.2.

6.7.4.2 Limited Aggregator availability

By removing the constraint that there are always as many Aggregators as Aggregation Ports, an implementation can limit the number of Aggregator Client interfaces available to higher layers while maintaining the ability for each Aggregator to serve multiple Aggregation Ports. This has the same effect as removing the assumption that Aggregators and their associated Aggregation Ports have the same operational Key value; Aggregators can be effectively disabled (and therefore ignored) by configuring their Keys to be different from any operational Key value allocated to any of the Aggregation Ports.

In this scenario, any Aggregation Port(s) that cannot find a suitable Aggregator to attach to will simply wait in the DETACHED state until an Aggregator becomes available, with a synchronization state of OUT_OF_SYNC.

7. Management

7.1 Overview

This clause provides the layer management specification for Link Aggregation. It provides the objects, attributes, and behaviours to support Link Aggregation.

The layout of this clause takes the same form as IEEE Std 802.3 Clause 30. It identifies a common management model and framework applicable to the IEEE 802.1AX managed elements, identifies those elements and defines their managed objects, attributes, and behaviours in a protocol-independent language.

It defines facilities comprised of a set of statistics and actions needed to provide IEEE 802.1AX management services. The Procedural Model provides a formal description of the relationship between Link Aggregation and the layer management facilities.

This management specification has been developed in accordance with the OSI management architecture as specified in the ISO Management Framework document, ISO/IEC 7498-4: 1989. It is independent of any particular management application or management protocol.

The management facilities defined in this standard may be accessed both locally and remotely. Thus, the layer management specification provides facilities that can be accessed from within a station or can be accessed remotely by means of a peer-management protocol operating between application entities.

In this standard, no peer management facilities are necessary for initiating or terminating normal protocol operations or for handling abnormal protocol conditions. Since these activities are subsumed by the normal operation of the protocol, they are not considered to be a function of layer management and are, therefore, not discussed in this clause.

Implementation of part or all of layer management is not a requirement for conformance to any other clause of this standard.

The improper use of some of the facilities described in this clause may cause serious disruption of the network. In accordance with ISO management architecture, any necessary security provisions should be provided by the agent in the Local System Environment. This can be in the form of specific security features or in the form of security features provided by the peer communication facilities.

7.1.1 Systems management overview

Within the ISO Open Systems Interconnection (OSI) architecture, the need to handle the special problems of initializing, terminating, and monitoring ongoing activities and assisting in their operations, as well as handling abnormal conditions, is recognized. These needs are collectively addressed by the systems management component of the OSI architecture.

A management protocol is required for the exchange of information between systems on a network. This management standard is independent of any particular management protocol.

This management standard, in conjunction with the management standards of other layers, provides the means to perform various management functions. IEEE 802.1AX management collects information needed from, and provides a means to exercise control over, Link Aggregation.

7.1.2 Management model

This standard describes management of Link Aggregation in terms of a general model of management of resources within the open systems environment. The model, which is described in ISO/IEC 10040: 1992, is briefly summarized here.

Management is viewed as a distributed application modeled as a set of interacting management processes. These processes are executed by systems within the open environment. A managing system executes a managing process that invokes management operations. A managed system executes a process that is receptive to these management operations and provides an interface to the resources to be managed. A managed object is the abstraction of a resource that represents its properties as seen by (and for the purpose of) management. Managed objects respond to a defined set of management operations. Managed objects are also capable of emitting a defined set of notifications.

A managed object is a management view of a resource. The resource may be a logical construct, function, physical device, or anything subject to management. Managed objects are defined in terms of four types of elements:

- a) *Attributes*. Data-like properties (as seen by management) of a managed object.
- b) *Actions*. Operations that a managing process may perform on an object or its attributes.
- c) *Notifications*. Unsolicited reports of events that may be generated by an object.
- d) *Behaviour*. The way in which managed objects, attributes, and actions interact with the actual resources they model and with each other.

The above items are defined in 7.3 of this clause in terms of the template requirements of ISO/IEC 10165-4: 1992.

Some of the functions and resources within IEEE 802.1AX devices are appropriate targets for management. They have been identified by specifying managed objects that provide a management view of the functions or resources. Within this general model, the IEEE 802.1AX device is viewed as a managed device. It performs functions as defined by the applicable standard for such a device. Managed objects providing a view of those functions and resources appropriate to the management of the device are specified. The purpose of this standard is to define the object classes associated with the devices in terms of their attributes, operations, notifications, and behaviour.

7.2 Managed objects

7.2.1 Introduction

This clause identifies the Managed Object classes for IEEE 802.1AX components within a managed system. It also identifies which managed objects and packages are applicable to which components.

All counters defined in this specification are assumed to be wrap-around counters. Wrap-around counters are those that automatically go from their maximum value (or final value) to zero and continue to operate. These unsigned counters do not provide for any explicit means to return them to their minimum (zero), i.e., reset. Because of their nature, wrap-around counters should be read frequently enough to avoid loss of information. When a counter has a maximum increment rate specified at one speed of operation, and that counter is appropriate to a higher speed of operation, then the maximum increment rate at that higher speed of operation is as shown in Equation (6-1).

$$\text{maximum increment rate specified} \times \left(\frac{\text{higher speed of operation in Mb/s}}{\text{specified speed of operation in Mb/s}} \right) \quad (6-1)$$

unless otherwise indicated.

7.2.2 Overview of managed objects

Managed objects provide a means to

- Identify a resource
- Control a resource
- Monitor a resource

7.2.2.1 Text description of managed objects

In case of conflict, the formal behaviour definitions in 7.3 take precedence over the text descriptions in this subclause.

oAggPortDebugInformation

A single instance of oAggPortDebugInformation may be contained within oAggregationPort. This managed object class provides optional additional information that can assist with debugging and fault finding in Systems that support Link Aggregation.

oAggPortStats

A single instance of oAggPortStats may be contained within oAggregationPort. This managed object class provides optional additional statistics related to LACP and Marker protocol activity on an instance of an Aggregation Port that is involved in Link Aggregation.

oAggregationPort

oAggregationPort is contained within oAggregator. An instance of this managed object class is present for each Aggregation Port that is part of the aggregation represented by the oAggregator instance. This managed object class provides the basic management controls necessary to allow an instance of an Aggregation Port to be managed, for the purposes of Link Aggregation.

oAggregator

oAggregator is contained within oDistributedRelay. Since oDistributedRelay is optional, oAggregator can be the top-most managed object class of the Link Aggregation containment tree. The oAggregator managed object class provides the management controls necessary to allow an instance of an Aggregator to be managed.

oDistributedRelay

oDistributedRelay is optional, and when present, is the top-most managed object class of the tree shown in Figure 7-1. Note that this managed object class (or oAggregator, if oDistributedRelay is not present) may be contained within another superior managed object class. Such containment is expected, but is outside the scope of this International Standard. The oDistributedRelay managed object class provides the management controls necessary to allow an instance of a Distributed Relay to be managed.

oDistributedRelayIPP

oDistributedRelayIPP is contained within oDistributedRelay. An instance of this managed object class is present for each IPP on the DR Function of a Portal System represented by the oDistributedRelay instance. This managed object class provides the basic management controls necessary to allow an instance of an IPP for the purposes of DRNI.

oIPPDebugInformation

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

A single instance of oIPPDebugInformation may be contained within oDistributedRelayIPP. This managed object class provides optional additional information that can assist with debugging and fault finding in Systems that support DRNI.

oIPPStats

A single instance of oIPPStats may be contained within oDistributedRelayIPP. This managed object class provides optional additional statistics related to DRCP protocol activity on an instance of an IPP that is involved in DRNI.

7.2.3 Containment

A containment relationship is a structuring relationship for managed objects in which the existence of a managed object is dependent on the existence of a containing managed object. The contained managed object is said to be the subordinate managed object, and the containing managed object the superior managed object. The containment relationship is used for naming managed objects. The local containment relationships among object classes are depicted in the entity relationship diagram Figure 7-1. The figure show the names of the object classes and whether a particular containment relationship is one-to-one or one-to-many. For further requirements on this topic, see IEEE Std 802.1F™-1993.

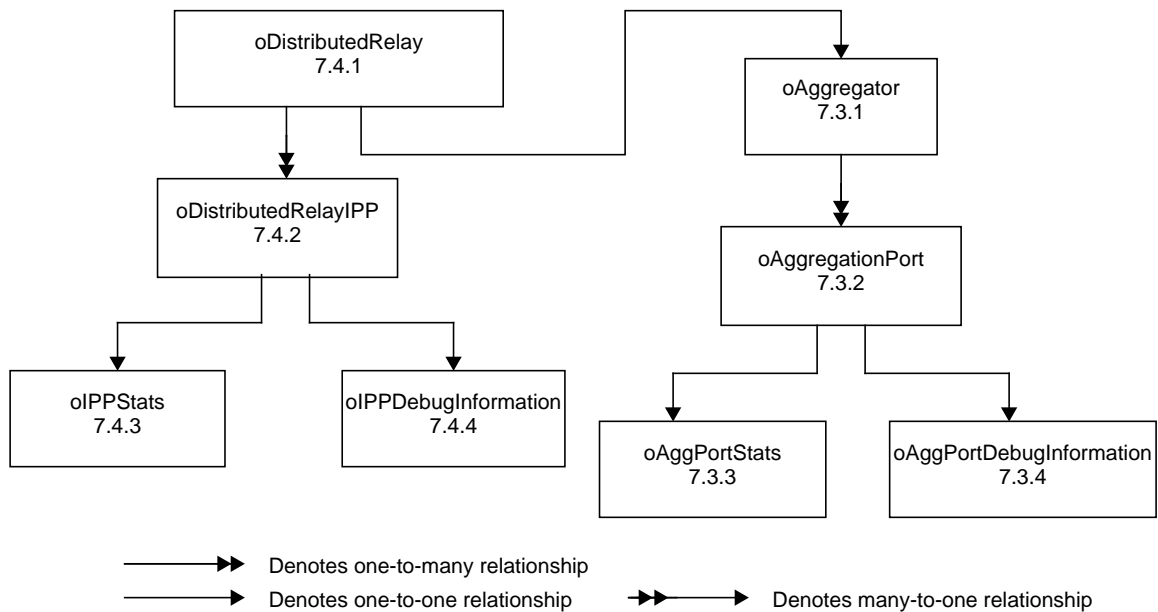


Figure 7-1—Link aggregation entity relationship diagram

7.2.4 Naming

The name of an individual managed object is hierarchically defined within a managed system. For example, an Aggregator might be identified as “aggregator 3, Aggregation Port 13” that is, Aggregation Port 13 of aggregator 3 within the managed system.

7.2.5 Capabilities

This standard makes use of the concept of *packages* as defined in ISO/IEC 10165-4: 1992 as a means of grouping behaviour, attributes, actions, and notifications within a managed object class definition. Packages

may either be mandatory, or be conditional, that is to say, present if a given condition is true. Within this standard *capabilities* are defined, each of which corresponds to a set of packages, which are components of a number of managed object class definitions and which share the same condition for presence. Implementation of the Basic and Mandatory packages is the minimum requirement for claiming conformance to IEEE 802.1AX Management. Implementation of an entire optional capability is required in order to claim conformance to that capability, unless stated otherwise. The capabilities and packages for IEEE 802.1AX Management are specified in Table 7–1.

Table 7–1—Link Aggregation capabilities

Object name	Object type	Operations supported	DTE									
			Basic package (mandatory)	Mandatory package (mandatory)	Recommended package (optional)	Optional package (optional)	Aggregation Port statistics (optional)	Aggregation Port debug information (optional)	Per-Service Frame Distribution package (optional)	DRNI package (optional)	IPP statistics (optional)	IPP debug information (optional)
oAggregator (7.3.1)												
aAggID	ATTRIBUTE	GET	X									
aAggDescription	ATTRIBUTE	GET	X									
aAggName	ATTRIBUTE	GET-SET	X									
aAggActorSystemID	ATTRIBUTE	GET-SET	X									
aAggActorSystemPriority	ATTRIBUTE	GET-SET	X									
aAggAggregateOrIndividual	ATTRIBUTE	GET	X									
aAggActorAdminKey	ATTRIBUTE	GET-SET	X									
aAggActorOperKey	ATTRIBUTE	GET	X									
aAggMACAddress	ATTRIBUTE	GET	X									
aAggPartnerSystemID	ATTRIBUTE	GET	X									
aAggPartnerSystemPriority	ATTRIBUTE	GET	X									
aAggPartnerOperKey	ATTRIBUTE	GET	X									
aAggAdminState	ATTRIBUTE	GET-SET	X									
aAggOperState	ATTRIBUTE	GET	X									
aAggTimeOfLastOperChange	ATTRIBUTE	GET	X									
aAggDataRate	ATTRIBUTE	GET	X									
aAggOctetsTxOK	ATTRIBUTE	GET		X								
aAggOctetsRxOK	ATTRIBUTE	GET		X								
aAggFramesTxOK	ATTRIBUTE	GET	X									
aAggFramesRxOK	ATTRIBUTE	GET	X									
aAggMulticastFramesTxOK	ATTRIBUTE	GET			X							
aAggMulticastFramesRxOK	ATTRIBUTE	GET			X							
aAggBroadcastFramesTxOK	ATTRIBUTE	GET			X							
aAggBroadcastFramesRxOK	ATTRIBUTE	GET			X							

Table 7–1—Link Aggregation capabilities (continued)

Object name	Object type	Operations supported	DTE									
			Basic package (mandatory)	Mandatory package (mandatory)	Recommended package (optional)	Optional package (optional)	Aggregation Port statistics (optional)	Aggregation Port debug information (optional)	Per-Service Frame Distribution package (optional)	DRNI package (optional)	IPP statistics (optional)	IPP debug information (optional)
aAggFramesDiscardedOnTx	ATTRIBUTE	GET		X								
aAggFramesDiscardedOnRx	ATTRIBUTE	GET		X								
aAggFramesWithTxErrors	ATTRIBUTE	GET		X								
aAggFramesWithRxErrors	ATTRIBUTE	GET		X								
aAggUnknownProtocolFrames	ATTRIBUTE	GET		X								
aAggLinkUpDownNotificationEnable	ATTRIBUTE	GET-SET	X									
nAggLinkUpNotification	NOTIFICATION		X									
nAggLinkDownNotification	NOTIFICATION		X									
aAggPortList	ATTRIBUTE	GET		X								
aAggCollectorMaxDelay	ATTRIBUTE	GET-SET	X									
aAggPortAlgorithm	ATTRIBUTE	GET-SET						X				
aAggPartnerAdminPortAlgorithm	ATTRIBUTE	GET-SET						X				
aAggConversationAdminLink[]	ATTRIBUTE	GET-SET						X				
aAggPartnerAdminPortConversationListDigest	ATTRIBUTE	GET-SET						X				
aAggAdminDiscardWrongConversation	ATTRIBUTE	GET-SET						X				
aAggAdminServiceConversationMap[]	ATTRIBUTE	GET-SET						X				
aAggPartnerAdminConvServiceMappingDigest	ATTRIBUTE	GET-SET						X				
oAggregationPort (7.3.2)												
aAggPortID	ATTRIBUTE	GET	X									
aAggPortActorSystemPriority	ATTRIBUTE	GET-SET		X								
aAggPortActorSystemID	ATTRIBUTE	GET		X								
aAggPortActorAdminKey	ATTRIBUTE	GET-SET		X								
aAggPortActorOperKey	ATTRIBUTE	GET		X								
aAggPortPartnerAdminSystemPriority	ATTRIBUTE	GET-SET		X								
aAggPortPartnerOperSystemPriority	ATTRIBUTE	GET		X								
aAggPortPartnerAdminSystemID	ATTRIBUTE	GET-SET		X								
aAggPortPartnerOperSystemID	ATTRIBUTE	GET		X								
aAggPortPartnerAdminKey	ATTRIBUTE	GET-SET		X								
aAggPortPartnerOperKey	ATTRIBUTE	GET		X								
aAggPortSelectedAggID	ATTRIBUTE	GET		X								
aAggPortAttachedAggID	ATTRIBUTE	GET		X								
aAggPortActorPort	ATTRIBUTE	GET		X								
aAggPortActorPortPriority	ATTRIBUTE	GET-SET		X								

Table 7–1—Link Aggregation capabilities (continued)

Object name	Object type	Operations supported	DTE									
			Basic package (mandatory)	Mandatory package (mandatory)	Recommended package (optional)	Optional package (optional)	Aggregation Port statistics (optional)	Aggregation Port debug information (optional)	Per-Service Frame Distribution package (optional)	DRNI package (optional)	IPP statistics (optional)	IPP debug information (optional)
aAggPortPartnerAdminPort	ATTRIBUTE	GET-SET	X									
aAggPortPartnerOperPort	ATTRIBUTE	GET	X									
aAggPortPartnerAdminPortPriority	ATTRIBUTE	GET-SET	X									
aAggPortPartnerOperPortPriority	ATTRIBUTE	GET	X									
aAggPortActorAdminState	ATTRIBUTE	GET-SET	X									
aAggPortActorOperState	ATTRIBUTE	GET	X									
aAggPortPartnerAdminState	ATTRIBUTE	GET-SET	X									
aAggPortPartnerOperState	ATTRIBUTE	GET	X									
aAggPortAggregateOrIndividual	ATTRIBUTE	GET	X									
aAggPortOperConversationPasses	ATTRIBUTE	GET							X			
aAggPortOperConversationCollected	ATTRIBUTE	GET							X			
aAggPortLinkNumberID	ATTRIBUTE	GET-SET							X			
aAggPortPartnerAdminLinkNumberID	ATTRIBUTE	GET-SET							X			
aAggPortWTRTime	ATTRIBUTE	GET-SET							X			
aAggPortProtocolDA	ATTRIBUTES	GET-SET	X									
oAggPortStats (7.3.3)												
aAggPortStatsID	ATTRIBUTE	GET						X				
aAggPortStatsLACPDUrx	ATTRIBUTE	GET						X				
aAggPortStatsMarkerPDUrx	ATTRIBUTE	GET						X				
aAggPortStatsMarkerResponsePDUrx	ATTRIBUTE	GET						X				
aAggPortStatsUnknownRx	ATTRIBUTE	GET						X				
aAggPortStatsIllegalRx	ATTRIBUTE	GET						X				
aAggPortStatsLACPDUtx	ATTRIBUTE	GET						X				
aAggPortStatsMarkerPDUtx	ATTRIBUTE	GET						X				
aAggPortStatsMarkerResponsePDUtx	ATTRIBUTE	GET						X				
oAggPortDebugInformation (7.3.4)												
aAggPortDebugInformationID	ATTRIBUTE	GET						X				
aAggPortDebugRxState	ATTRIBUTE	GET						X				
aAggPortDebugLastRxTime	ATTRIBUTE	GET						X				
aAggPortDebugMuxState	ATTRIBUTE	GET						X				
aAggPortDebugMuxReason	ATTRIBUTE	GET						X				
aAggPortDebugActorChurnState	ATTRIBUTE	GET						X				
aAggPortDebugPartnerChurnState	ATTRIBUTE	GET						X				

Table 7–1—Link Aggregation capabilities (continued)

Object name	Object type	Operations supported	DTE										
			Basic package (mandatory)	Mandatory package (mandatory)	Recommended package (optional)	Optional package (optional)	Aggregation Port statistics (optional)	Aggregation Port debug information (optional)	Per-Service Frame Distribution package (optional)	DRNI package (optional)	IPP statistics (optional)	IPP debug information (optional)	
aAggPortDebugActorChurnCount	ATTRIBUTE	GET							X				
aAggPortDebugPartnerChurnCount	ATTRIBUTE	GET							X				
aAggPortDebugActorSyncTransitionCount	ATTRIBUTE	GET							X				
aAggPortDebugPartnerSyncTransitionCount	ATTRIBUTE	GET							X				
aAggPortDebugActorChangeCount	ATTRIBUTE	GET							X				
aAggPortDebugPartnerChangeCount	ATTRIBUTE	GET							X				
aAggPortDebugActorCDSChurnState	ATTRIBUTE	GET							X				
aAggPortDebugPartnerCDSChurnState	ATTRIBUTE	GET							X				
aAggPortDebugActorCDSChurnCount	ATTRIBUTE	GET							X				
aAggPortDebugPartnerCDSChurnCount	ATTRIBUTE	GET							X				
oDistributedRelay (7.4.1)													
aDrniID	ATTRIBUTE	GET										X	
aDrniDescription	ATTRIBUTE	GET										X	
aDrniName	ATTRIBUTE	GET-SET										X	
aDrniPortalAddr	ATTRIBUTE	GET-SET										X	
aDrniPortalPriority	ATTRIBUTE	GET-SET										X	
aDrniThreePortalSystem	ATTRIBUTE	GET-SET										X	
aDrniPortalSystemNumber	ATTRIBUTE	GET-SET										X	
aDrniIntraPortalLinkList	ATTRIBUTE	GET-SET										X	
aDrniAggregator	ATTRIBUTE	GET-SET										X	
aDrniConvAdminGateway[]	ATTRIBUTE	GET-SET										X	
aDrniNeighborAdminConvGatewayListDigest	ATTRIBUTE	GET-SET										X	
aDrniNeighborAdminConvPortListDigest	ATTRIBUTE	GET-SET										X	
aDrniGatewayAlgorithm	ATTRIBUTE	GET-SET										X	
aDrniNeighborAdminGatewayAlgorithm	ATTRIBUTE	GET-SET										X	
aDrniNeighborAdminPortAlgorithm	ATTRIBUTE	GET-SET										X	
aDrniNeighborAdminDRCPState	ATTRIBUTE	GET-SET										X	
aDrniEncapsulationMethod	ATTRIBUTE	GET-SET										X	
aDrniPLEncapMap	ATTRIBUTE	GET-SET										X	
aDrniNetEncapMap	ATTRIBUTE	GET-SET										X	
aDrniDRPortConversationPasses	ATTRIBUTE	GET										X	
aDrniDRGatewayConversationPasses	ATTRIBUTE	GET										X	
aDrniPSI	ATTRIBUTE	GET										X	

Table 7–1—Link Aggregation capabilities (continued)

Object name	Object type	Operations supported	DTE											
			Basic package (mandatory)	Mandatory package (mandatory)	Recommended package (optional)	Optional package (optional)	Aggregation Port statistics (optional)	Aggregation Port debug information (optional)	Per-Service Frame Distribution package (optional)	DRNI package (optional)	IPP statistics (optional)	IPP debug information (optional)		
aDrniPortConversationControl	ATTRIBUTES	GET-SET										X		
aDrniIntraPortalPortProtocolDA	ATTRIBUTES	GET-SET										X		
oDistributedRelayIPP (7.4.2)														
aIPPID	ATTRIBUTE	GET										X		
aIPPPortConversationPasses	ATTRIBUTE	GET										X		
aIPPGatewayConversationDirection	ATTRIBUTE	GET										X		
aIPPAdminState	ATTRIBUTE	GET-SET										X		
aIPPOperState	ATTRIBUTE	GET										X		
aIPPTimeOfLastOperChange	ATTRIBUTE	GET										X		
oIPPStats (7.4.3)														
aIPPStatsID	ATTRIBUTE	GET											X	
aIPPStatsDRCPDUsRx	ATTRIBUTE	GET											X	
aIPPStatsIllegalRx	ATTRIBUTE	GET											X	
aIPPStatsDRCPDUsTx	ATTRIBUTE	GET											X	
oIPPDebugInformation (7.4.4)														
aIPPDebugInformationID	ATTRIBUTE	GET												X
aIPPDebugDRCPRxState	ATTRIBUTE	GET												X
aIPPDebugLastRxTime	ATTRIBUTE	GET												X
aIPPDebugDifferPortalReason	ATTRIBUTE	GET												X
Common Attributes Template														
aCMCounter	ATTRIBUTE	GET	X	X	X	X	X	X					X	X

7.3 Management for Link Aggregation

7.3.1 Aggregator managed object class

This subclause formally defines the behaviours for the oAggregator managed object class, attributes, and notifications.

Some of the attributes that are part of the definition of the oAggregator managed object class are derived by summing counter values from attributes of other objects; e.g., to generate a count of received frames for the Aggregator, the individual value for each Aggregation Port contributes to the sum. Where calculations of

1 this form are used, the values that contribute to the Aggregator's attributes are *increments* in the values of
2 the component attributes, not their absolute values. As any individual Aggregation Port is potentially only
3 temporarily attached to its current Aggregator, the count values it contributes to the Aggregator's counters
4 are the increments in its values that it has experienced during the period of time that it has been attached to
5 that Aggregator.
6

7 The counter values defined for the Aggregator have been formulated as far as possible to make the
8 Aggregator behave like an individual IEEE 802 MAC. The counts of frames received and transmitted are
9 formulated to reflect the counts that would be expected by the Aggregator Client; they do not include frames
10 transmitted and received as part of the operation of LACP or the Marker protocol, only frames that pass
11 through the interface between the Aggregator Client and the Aggregator. However, as LACP and the Marker
12 protocol are, as far as the individual MACs are concerned, part of their Aggregator Client, the RX/TX
13 counters for the individual MACs will reflect both control and data traffic. As counts of errors at the
14 Aggregation Port level cannot always be cleanly delineated between those that occurred as a result of
15 aggregation activity and those that did not, no attempt has been made to separate these aspects of the
16 Aggregation Port error counts. Therefore, there is not necessarily a direct correspondence between the
17 individual MAC counters and the corresponding derived counters at the Aggregator level.
18

19 It should also be noted that the counters defined for the Aggregator include values that can only apply to half
20 duplex links. This is consistent with the approach taken in Link Aggregation that a link that can only operate
21 as an individual link is nonetheless considered as being attached to an Aggregator. This simplifies the
22 modelling of managed objects for links that can operate in either half or full duplex, and ensures a consistent
23 presentation of the attributes regardless of the type of links attached to the Aggregator.
24

25 NOTE—The operation of Auto-Negotiation may mean that a given link can operate in full duplex or half duplex,
26 depending upon the capabilities of the device(s) connected to it. Keeping the management view the same regardless of a
27 link's current mode of operation allows a consistent management approach to be taken across all types of links.
28

29 **7.3.1.1 Aggregator attributes**

30 **7.3.1.1.1 aAggID**

31 ATTRIBUTE

32 APPROPRIATE SYNTAX:

33 INTEGER

34 BEHAVIOUR DEFINED AS:

35
36 The unique identifier allocated to this Aggregator by the local System. This attribute
37 identifies an Aggregator instance among the subordinate managed objects of the containing
38 object. This value is read-only.
39
40
41

42 NOTE—The aAggID is represented in the SMIv2 MIB as an ifIndex—see D.4.1.
43

44 **7.3.1.1.2 aAggDescription**

45 ATTRIBUTE

46 APPROPRIATE SYNTAX:

47 A PrintableString, 255 characters max.

48 BEHAVIOUR DEFINED AS:

49 A human-readable text string containing information about the Aggregator. This string could
50 include information about the distribution algorithm in use on this Aggregator; for example,
51 "Aggregator 1, Dist Alg=Dest MAC address." This string is read-only. The contents are
52 vendor specific.
53
54

7.3.1.1.3 aAggName

ATTRIBUTE

APPROPRIATE SYNTAX:

A PrintableString, 255 characters max.

BEHAVIOUR DEFINED AS:

A human-readable text string containing a locally significant name for the Aggregator. This string is read-write.

7.3.1.1.4 aAggActorSystemID

ATTRIBUTE

APPROPRIATE SYNTAX:

MACAddress

BEHAVIOUR DEFINED AS:

A 6-octet read-write MAC address value used as a unique identifier for the System that contains this Aggregator.

NOTE—From the perspective of the Link Aggregation mechanisms described in Clause 6, only a single combination of Actor's System ID and System Priority are considered, and no distinction is made between the values of these parameters for an Aggregator and the Aggregation Port(s) that are associated with it (i.e., the protocol is described in terms of the operation of aggregation within a single System). However, the managed objects provided for the Aggregator and the Aggregation Port both allow management of these parameters. The result of this is to permit a single piece of equipment to be configured by management to contain more than one System from the point of view of the operation of Link Aggregation. This may be of particular use in the configuration of equipment that has limited aggregation capability (see 6.7).

7.3.1.1.5 aAggActorSystemPriority

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

A 2-octet read-write value indicating the priority value associated with the Actor's System ID.

7.3.1.1.6 aAggAggregateOrIndividual

ATTRIBUTE

APPROPRIATE SYNTAX:

BOOLEAN

BEHAVIOUR DEFINED AS:

A read-only Boolean value indicating whether the Aggregator represents an Aggregate ("TRUE") or an Individual link ("FALSE").

7.3.1.1.7 aAggActorAdminKey

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The current administrative value of the Key for the Aggregator. The administrative Key value may differ from the operational Key value for the reasons discussed in 6.7.2. This is a 16-bit

1 read-write value. The meaning of particular Key values is of local significance. For an
2 Aggregator that is associated with a Portal the aAggActorAdminKey has to be different for
3 each Portal System. Specifically the two most significant bits are set to
4 aDmniPortalSystemNumber (7.4.1.1.7). The lower 14 bits may be any value, have to be the
5 same in each Portal System within the same Portal, and have a default of zero.

7.3.1.1.8 aAggActorOperKey

8 ATTRIBUTE

9 APPROPRIATE SYNTAX:

10 INTEGER

11 BEHAVIOUR DEFINED AS:

12 The current operational value of the Key for the Aggregator. The administrative Key value
13 may differ from the operational Key value for the reasons discussed in 6.7.2. This is a 16-bit
14 read-only value. The meaning of particular Key values is of local significance.

7.3.1.1.9 aAggMACAddress

15 ATTRIBUTE

16 APPROPRIATE SYNTAX:

17 MACAddress

18 BEHAVIOUR DEFINED AS:

19 A 6-octet read-only value carrying the individual MAC address assigned to the Aggregator.

7.3.1.1.10 aAggPartnerSystemID

20 ATTRIBUTE

21 APPROPRIATE SYNTAX:

22 MACAddress

23 BEHAVIOUR DEFINED AS:

24 A 6-octet read-only MAC address value consisting of the unique identifier for the current
25 protocol Partner of this Aggregator. A value of zero indicates that there is no known Partner.
26 If the aggregation is manually configured, this System ID value will be a value assigned by
27 the local System.

7.3.1.1.11 aAggPartnerSystemPriority

28 ATTRIBUTE

29 APPROPRIATE SYNTAX:

30 INTEGER

31 BEHAVIOUR DEFINED AS:

32 A 2-octet read-only value that indicates the priority value associated with the Partner's
33 System ID. If the aggregation is manually configured, this System Priority value will be a
34 value assigned by the local System.

7.3.1.1.12 aAggPartnerOperKey

35 ATTRIBUTE

36 APPROPRIATE SYNTAX:

37 INTEGER

38 BEHAVIOUR DEFINED AS:

1 The current operational value of the Key for the Aggregator’s current protocol Partner. This
2 is a 16-bit read-only value. If the aggregation is manually configured, this Key value will be a
3 value assigned by the local System.
4

5 **7.3.1.1.13 aAggAdminState**

6 ATTRIBUTE

7 APPROPRIATE SYNTAX:

8 An ENUMERATED VALUE that has one of the following entries:

9 up
10 down

11 BEHAVIOUR DEFINED AS:

12 This read-write value defines the administrative state of the Aggregator. A value of “up”
13 indicates that the operational state of the Aggregator (aAggOperState) is permitted to be
14 either up or down. A value of “down” forces the operational state of the Aggregator to be
15 down. Changes to the administrative state affect the operational state of the Aggregator only,
16 not the operational state of the Aggregation Ports that are attached to the Aggregator. A GET
17 operation returns the current administrative state. A SET operation changes the
18 administrative state to a new value.
19

20 **7.3.1.1.14 aAggOperState**

21 ATTRIBUTE

22 APPROPRIATE SYNTAX:

23 An ENUMERATED VALUE that has one of the following entries:

24 up
25 down

26 BEHAVIOUR DEFINED AS:

27 This read-only value defines the operational state of the Aggregator. An operational state of
28 “up” indicates that the Aggregator is available for use by the Aggregator Client; a value of
29 “down” indicates that the Aggregator is not available for use by the Aggregator Client.
30

31 **7.3.1.1.15 aAggTimeOfLastOperChange**

32 ATTRIBUTE

33 APPROPRIATE SYNTAX:

34 INTEGER

35 BEHAVIOUR DEFINED AS:

36 The time at which the interface entered its current operational state, in terms of centiseconds
37 since the system was last reset. If the current state was entered prior to the last re-
38 initialization of the local network management subsystem, then this object contains a value of
39 zero. The ifLastChange object in the Interfaces MIB defined in IETF RFC 2863 is a suitable
40 object for supplying a value for aAggTimeOfLastOperChange. This value is read-only.
41

42 NOTE—aAggTimeOfLastOperChange was defined in terms of the aTimeSinceSystemReset variable of IEEE Std 802.3
43 Annex F.2.1 in earlier versions of this standard. aTimeSinceSystemReset and ifLastChange have the meaning.
44

45 **7.3.1.1.16 aAggDataRate**

46 ATTRIBUTE

47 APPROPRIATE SYNTAX:

48 INTEGER
49

BEHAVIOUR DEFINED AS:

The current data rate, in bits per second, of the aggregate link. The value is calculated as the sum of the data rate of each link in the aggregation. This attribute is read-only.

7.3.1.1.17 aAggOctetsTxOK**ATTRIBUTE****APPROPRIATE SYNTAX:**

Generalized counter. This counter has a maximum increment rate of 1 230 000 counts per second for a single 10 Mb/s aggregation.

BEHAVIOUR DEFINED AS:

A count of the data and padding octets transmitted by this Aggregator on all Aggregation Ports that are (or have been) members of the aggregation. The count does not include octets transmitted by the Aggregator in frames that carry LACPDUs or Marker PDUs (7.3.3.1.7, 7.3.3.1.8, 7.3.3.1.9). However, it includes frames discarded by the Frame Distribution function of the Aggregator (7.3.1.1.25). This value is read-only.

7.3.1.1.18 aAggOctetsRxOK**ATTRIBUTE****APPROPRIATE SYNTAX:**

Generalized counter. This counter has a maximum increment rate of 1 230 000 counts per second for a single 10 Mb/s aggregation.

BEHAVIOUR DEFINED AS:

A count of the data and padding octets received by this Aggregator, from the Aggregation Ports that are (or have been) members of the aggregation. The count does not include octets received in frames that carry LACP or Marker PDUs (7.3.3.1.2, 7.3.3.1.3, 7.3.3.1.4), or frames discarded by the Frame Collection function of the Aggregator (7.3.1.1.26). This value is read-only.

7.3.1.1.19 aAggFramesTxOK**ATTRIBUTE****APPROPRIATE SYNTAX:**

Generalized counter. This counter has a maximum increment rate of 16 000 counts per second for a single 10 Mb/s aggregation.

BEHAVIOUR DEFINED AS:

A count of the data frames transmitted by this Aggregator on all Aggregation Ports that are (or have been) members of the aggregation. The count does not include frames transmitted by the Aggregator that carry LACP or Marker PDUs (7.3.3.1.7, 7.3.3.1.8, 7.3.3.1.9). However, it includes frames discarded by the Frame Distribution function of the Aggregator (7.3.1.1.25). This value is read-only.

7.3.1.1.20 aAggFramesRxOK**ATTRIBUTE****APPROPRIATE SYNTAX:**

Generalized counter. This counter has a maximum increment rate of 16 000 counts per second for a single 10 Mb/s aggregation.

BEHAVIOUR DEFINED AS:

A count of the data frames received by this Aggregator, from the Aggregation Ports that are (or have been) members of the aggregation. The count does not include frames that carry

1 LACP or Marker PDUs (7.3.3.1.2, 7.3.3.1.3, 7.3.3.1.4), or frames discarded by the Frame
2 Collection function of the Aggregator (7.3.1.1.26). This value is read-only.

3 4 **7.3.1.1.21 aAggMulticastFramesTxOK**

5
6 ATTRIBUTE

7 APPROPRIATE SYNTAX:

8 Generalized counter. This counter has a maximum increment rate of 16 000 counts per
9 second for a single 10 Mb/s aggregation.

10 BEHAVIOUR DEFINED AS:

11 A count of the data frames transmitted by this Aggregator on all Aggregation Ports that are
12 (or have been) members of the aggregation, to a group destination address other than the
13 broadcast address. The count does not include frames transmitted by the Aggregator that
14 carry LACP or Marker PDUs (7.3.3.1.7, 7.3.3.1.8, 7.3.3.1.9). However, it includes frames
15 discarded by the Frame Distribution function of the Aggregator (7.3.1.1.25). This value is
16 read-only.
17

18 **7.3.1.1.22 aAggMulticastFramesRxOK**

19
20 ATTRIBUTE

21 APPROPRIATE SYNTAX:

22 Generalized counter. This counter has a maximum increment rate of 16 000 counts per
23 second for a single 10 Mb/s aggregation.

24 BEHAVIOUR DEFINED AS:

25 A count of the data frames received by this Aggregator, from the Aggregation Ports that are
26 (or have been) members of the aggregation, that were addressed to an active group address
27 other than the broadcast address. The count does not include frames that carry LACP or
28 Marker PDUs (7.3.3.1.2, 7.3.3.1.3, 7.3.3.1.4), or frames discarded by the Frame Collection
29 function of the Aggregator (7.3.1.1.26). This value is read-only.
30
31

32 **7.3.1.1.23 aAggBroadcastFramesTxOK**

33
34 ATTRIBUTE

35 APPROPRIATE SYNTAX:

36 Generalized counter. This counter has a maximum increment rate of 16 000 counts per
37 second for a single 10 Mb/s aggregation.

38 BEHAVIOUR DEFINED AS:

39 A count of the broadcast data frames transmitted by this Aggregator on all Aggregation Ports
40 that are (or have been) members of the aggregation. The count does not include frames
41 transmitted by the Aggregator that carry LACP or Marker PDUs (7.3.3.1.7, 7.3.3.1.8,
42 7.3.3.1.9). However, it includes frames discarded by the Frame Distribution function of the
43 Aggregator (7.3.1.1.25). This value is read-only.
44

45 **7.3.1.1.24 aAggBroadcastFramesRxOK**

46
47 ATTRIBUTE

48 APPROPRIATE SYNTAX:

49 Generalized counter. This counter has a maximum increment rate of 16 000 counts per
50 second for a single 10 Mb/s aggregation.

51 BEHAVIOUR DEFINED AS:

52 A count of the broadcast data frames received by this Aggregator, from the Aggregation Ports
53 that are (or have been) members of the aggregation. The count does not include frames that
54

1 carry LACP or Marker PDUs (7.3.3.1.2, 7.3.3.1.3, 7.3.3.1.4), illegal or unknown protocol
2 frames (7.3.3.1.5, 7.3.3.1.6), or frames discarded by the Frame Collection function of the
3 Aggregator (7.3.1.1.26). This value is read-only.
4

5 **7.3.1.1.25 aAggFramesDiscardedOnTx**

6 ATTRIBUTE

7 APPROPRIATE SYNTAX:

8 Generalized counter. This counter has a maximum increment rate of 16 000 counts per
9 second for a single 10 Mb/s aggregation.
10

11 BEHAVIOUR DEFINED AS:

12 A count of data frames requested to be transmitted by this Aggregator that were discarded by
13 the Frame Distribution function of the Aggregator when conversations are re-allocated to
14 different Aggregation Ports, due to the requirement to ensure that the conversations are
15 flushed on the old Aggregation Ports in order to maintain proper frame ordering (43A.3), or
16 discarded as a result of excessive collisions by Aggregation Ports that are (or have been)
17 members of the aggregation. This value is read-only.
18

19 **7.3.1.1.26 aAggFramesDiscardedOnRx**

20 ATTRIBUTE

21 APPROPRIATE SYNTAX:

22 Generalized counter. This counter has a maximum increment rate of 16 000 counts per
23 second for a single 10 Mb/s aggregation.
24

25 BEHAVIOUR DEFINED AS:

26 A count of data frames, received on all Aggregation Ports that are (or have been) members of
27 the aggregation, that were discarded by the Frame Collection function of the Aggregator as
28 they were received on Aggregation Ports whose Frame Collection function was disabled.
29 This value is read-only.
30

31 **7.3.1.1.27 aAggFramesWithTxErrors**

32 ATTRIBUTE

33 APPROPRIATE SYNTAX:

34 Generalized counter. This counter has a maximum increment rate of 16 000 counts per
35 second for a single 10 Mb/s aggregation.
36

37 BEHAVIOUR DEFINED AS:

38 A count of data frames requested to be transmitted by this Aggregator that experienced
39 transmission errors on Aggregation Ports that are (or have been) members of the aggregation.
40 This count does not include frames discarded due to excess collisions. This value is read-
41 only.
42
43

44 **7.3.1.1.28 aAggFramesWithRxErrors**

45 ATTRIBUTE

46 APPROPRIATE SYNTAX:

47 Generalized counter. This counter has a maximum increment rate of 16 000 counts per
48 second for a single 10 Mb/s aggregation.
49

50 BEHAVIOUR DEFINED AS:

51 A count of data frames discarded on reception by all Aggregation Ports that are (or have
52 been) members of the aggregation, or that were discarded by the Frame Collection function
53
54

1 of the Aggregator, or that were discarded by the Aggregator due to the detection of an illegal
2 Slow Protocols PDU (7.3.3.1.6). This value is read-only.

3 4 **7.3.1.1.29 aAggUnknownProtocolFrames**

5
6 ATTRIBUTE

7 APPROPRIATE SYNTAX:

8 Generalized counter. This counter has a maximum increment rate of 16 000 counts per
9 second for a single 10 Mb/s aggregation.

10 BEHAVIOUR DEFINED AS:

11 A count of data frames discarded on reception by all Aggregation Ports that are (or have
12 been) members of the aggregation, due to the detection of an unknown Slow Protocols PDU
13 (7.3.3.1.5). This value is read-only.
14

15 16 **7.3.1.1.30 aAggPortList**

17 ATTRIBUTE

18 APPROPRIATE SYNTAX:

19 A SEQUENCE OF INTEGERS that match the syntax of aAggPortID.

20 BEHAVIOUR DEFINED AS:

21 The value of this read-only attribute contains the list of Aggregation Ports that are currently
22 attached to the Aggregator (6.3.9, 6.3.14, 6.4.15). An empty list indicates that there are no
23 Aggregation Ports attached. Each integer value in the list carries an aAggPortID attribute
24 value (7.3.2.1.1).
25
26

27 28 **7.3.1.1.31 aAggLinkUpDownNotificationEnable**

29 ATTRIBUTE

30 APPROPRIATE SYNTAX:

31 An ENUMERATED VALUE that has one of the following entries:

32 enabled

33 disabled

34 BEHAVIOUR DEFINED AS:

35 When set to “enabled,” Link Up and Link Down notifications are enabled for this
36 Aggregator. When set to “disabled,” Link Up and Link Down notifications are disabled for
37 this Aggregator. This value is read-write.
38

39 40 **7.3.1.1.32 aAggCollectorMaxDelay**

41 ATTRIBUTE

42 APPROPRIATE SYNTAX:

43 INTEGER

44 BEHAVIOUR DEFINED AS:

45 The value of this 16-bit read-write attribute defines the maximum delay, in tens of
46 microseconds, that may be imposed by the Frame Collector between receiving a frame from
47 an Aggregator Parser, and either delivering the frame to its Aggregator Client or discarding
48 the frame (see 6.2.3.1.1).
49
50

51 52 **7.3.1.1.33 aAggPortAlgorithm**

53 ATTRIBUTE
54

1 APPROPRIATE SYNTAX

2 A SEQUENCE OF OCTETS consisting of a three-octet OUI or CID and one following octet.

3 BEHAVIOR DEFINED AS

4 This object identifies the algorithm used by the Aggregator to assign frames to a Port
5 Conversation ID. Table 6-4 provides the IEEE 802.1 OUI (00-80-C2) Port Algorithm
6 encodings.
78 **7.3.1.1.34 aAggPartnerAdminPortAlgorithm**

9 ATTRIBUTE

10 APPROPRIATE SYNTAX

11 A SEQUENCE OF OCTETS consisting of a three-octet OUI or CID and one following octet.

12 BEHAVIOR DEFINED AS

13 This object identifies the value for the algorithm of the Partner System, assigned by
14 administrator or System policy for use when the Partner's information is unknown. Table 6-4
15 provides the IEEE 802.1 OUI (00-80-C2) Port Algorithm encodings. Its default value is set to
16 NULL.
17
1819 **7.3.1.1.35 aAggConversationAdminLink[]**

20 ATTRIBUTE

21 APPROPRIATE SYNTAX

22 An array of SEQUENCE OF INTEGERS that match the syntax of a Link Number ID
23 (6.6.2.2).
24

25 BEHAVIOR DEFINED AS

26 There are 4096 aAggConversationAdminLink[] variables, aAggConversationAdminLink[0]
27 through aAggConversationAdminLink[4095], indexed by Port Conversation ID. Each
28 contains administrative values of the link selection priority list for the referenced Port
29 Conversation ID. This selection priority list is a sequence of Link Number IDs for each Port
30 Conversation ID, in the order of preference, highest to lowest, for the corresponding link to
31 carry that Port Conversation ID. A 16 bit zero value is used to indicate that no link is assigned
32 to carry the associated Port Conversation ID.
33
3435 NOTE 1—This mapping of Port Conversation IDs to Link Number IDs is the fundamental administrative input. An
36 equivalent mapping of Port Conversation IDs to Port IDs [Conversation_PortList[] (6.6.2.1)] is derived from this and
37 used internally.
3839 NOTE 2—When a network administrator issues a command for selection rules, provided by
40 aAggConversationAdminLink[], and accompanied with a non-zero value for aAggPortWTRTime (7.3.2.1.29) for all
41 associated Aggregation Ports, the ChangeActorOperDist is set as specified in 6.6.2.2. A value of 100 for the
42 aAggPortWTRTime indicates a non-revertive mode of operation and the WTR_timer will be kept to the value 100.
4344 **7.3.1.1.36 aAggPartnerAdminPortConversationListDigest**

45 ATTRIBUTE

46 APPROPRIATE SYNTAX

47 A SEQUENCE OF OCTETS consisting of a 16-octet MD5 Digest.

48 BEHAVIOR DEFINED AS

49 The value for the digest of the prioritized Port Conversation ID-to-Link Number ID
50 assignments of the Partner System, assigned by administrator or System policy for use when
51 the Partner's information is unknown. Its default value is set to NULL.
52
53
54

7.3.1.1.37 aAggAdminDiscardWrongConversation

ATTRIBUTE

APPROPRIATE SYNTAX

BOOLEAN

BEHAVIOR DEFINED AS

The administrative value that determines what the Aggregator does with a frame that is received from an Aggregation Port with a Port Conversation ID that is not included in the Collection_Conversation_Mask. The value “TRUE” indicates that such frames are to be discarded, and the value “FALSE” that they are to be forwarded. This variable needs to be set to “TRUE”, if bidirectional congruity (8.2.1) is required. Its value is set to “TRUE” by default.

7.3.1.1.38 aAggAdminServiceConversationMap[]

ATTRIBUTE

APPROPRIATE SYNTAX

An array of SEQUENCE OF INTEGERS, that match the syntax of Service IDs (8.2.2).

BEHAVIOR DEFINED AS

There are 4096 aAggAdminServiceConversationMap[] variables, aAggAdminServiceConversationMap[0] through aAggAdminServiceConversationMap[4095], indexed by Port Conversation ID. Each contains, in general, a set of Service IDs (8.2.2), unique within the array. If the Service IDs are representing VIDs, only a single VID is used, while in the case that Service IDs are representing I-SIDs, more than one I-SIDs are possible. Service IDs not contained in the map are not mapped to any Port Conversation ID and will be discarded.

7.3.1.1.39 aAggPartnerAdminConvServiceMappingDigest

ATTRIBUTE

APPROPRIATE SYNTAX

A SEQUENCE OF OCTETS consisting of a 16-octet MD5 Digest.

BEHAVIOR DEFINED AS

The value for the digest of the Port Conversation ID-to-Service ID assignments of the Partner System, assigned by administrator or System policy for use when the Partner’s information is unknown. Its default value is set to NULL.

7.3.1.2 Aggregator Notifications

7.3.1.2.1 nAggLinkUpNotification

NOTIFICATION

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

When aAggLinkUpDownNotificationEnable is set to “enabled,” a Link Up notification is generated when the Operational State of the Aggregator changes from “down” to “up.” When aAggLinkUpDownNotificationEnable is set to “disabled,” no Link Up notifications are generated. The notification carries the identifier of the Aggregator whose state has changed.

7.3.1.2.2 nAggLinkDownNotification

NOTIFICATION

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

When aAggLinkUpDownNotificationEnable is set to “enabled,” a Link Down notification is generated when the Operational State of the Aggregator changes from “up” to “down.” When aAggLinkUpDownNotificationEnable is set to “disabled,” no Link Down notifications are generated. The notification carries the identifier of the Aggregator whose state has changed.

7.3.2 Aggregation Port managed object class

This subclause formally defines the behaviours for the oAggregationPort managed object class attributes.

7.3.2.1 Aggregation Port Attributes

7.3.2.1.1 aAggPortID

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The unique identifier allocated to this Aggregation Port by the local System. This attribute identifies an Aggregation Port instance among the subordinate managed objects of the containing object. This value is read-only.

NOTE—The aAggPortID is represented in the SMIV2 MIB as an ifIndex—see D.4.1.

7.3.2.1.2 aAggPortActorSystemPriority

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

A 2-octet read-write value used to define the priority value associated with the Actor’s System ID.

7.3.2.1.3 aAggPortActorSystemID

ATTRIBUTE

APPROPRIATE SYNTAX:

MACAddress

BEHAVIOUR DEFINED AS:

A 6-octet read-only MAC address value that defines the value of the System ID for the System that contains this Aggregation Port.

7.3.2.1.4 aAggPortActorAdminKey

ATTRIBUTE

APPROPRIATE SYNTAX:

1 INTEGER

2 BEHAVIOUR DEFINED AS:

3 The current administrative value of the Key for the Aggregation Port. This is a 16-bit read-
4 write value. The meaning of particular Key values is of local significance.
5

6 **7.3.2.1.5 aAggPortActorOperKey**

7
8 ATTRIBUTE

9 APPROPRIATE SYNTAX:

10 INTEGER

11 BEHAVIOUR DEFINED AS:

12 The current operational value of the Key for the Aggregation Port. This is a 16-bit read-only
13 value. The meaning of particular Key values is of local significance.
14

15 **7.3.2.1.6 aAggPortPartnerAdminSystemPriority**

16
17 ATTRIBUTE

18 APPROPRIATE SYNTAX:

19 INTEGER

20 BEHAVIOUR DEFINED AS:

21 A 2-octet read-write value used to define the administrative value of priority associated with
22 the Partner's System ID. The assigned value is used, along with the value of
23 aAggPortPartnerAdminSystemID, aAggPortPartnerAdminKey, aAggPortPartnerAdminPort,
24 and aAggPortPartnerAdminPortPriority, in order to achieve manually configured
25 aggregation.
26
27

28 **7.3.2.1.7 aAggPortPartnerOperSystemPriority**

29
30 ATTRIBUTE

31 APPROPRIATE SYNTAX:

32 INTEGER

33 BEHAVIOUR DEFINED AS:

34 A 2-octet read-only value indicating the operational value of priority associated with the
35 Partner's System ID. The value of this attribute may contain the manually configured value
36 carried in aAggPortPartnerAdminSystemPriority if there is no protocol Partner.
37
38

39 **7.3.2.1.8 aAggPortPartnerAdminSystemID**

40
41 ATTRIBUTE

42 APPROPRIATE SYNTAX:

43 MACAddress

44 BEHAVIOUR DEFINED AS:

45 A 6-octet read-write MACAddress value representing the administrative value of the
46 Aggregation Port's protocol Partner's System ID. The assigned value is used, along with the
47 value of aAggPortPartnerAdminSystemPriority, aAggPortPartnerAdminKey,
48 aAggPortPartnerAdminPort, and aAggPortPartnerAdminPortPriority, in order to achieve
49 manually configured aggregation.
50

51 **7.3.2.1.9 aAggPortPartnerOperSystemID**

52
53 ATTRIBUTE
54

1 APPROPRIATE SYNTAX:

2 MACAddress

3 BEHAVIOUR DEFINED AS:

4 A 6-octet read-only MACAddress value representing the current value of the Aggregation
5 Port's protocol Partner's System ID. A value of zero indicates that there is no known protocol
6 Partner. The value of this attribute may contain the manually configured value carried in
7 aAggPortPartnerAdminSystemID if there is no protocol Partner.
8

9 **7.3.2.1.10 aAggPortPartnerAdminKey**

10 ATTRIBUTE

11 APPROPRIATE SYNTAX:

12 INTEGER

13 BEHAVIOUR DEFINED AS:

14 The current administrative value of the Key for the protocol Partner. This is a 16-bit
15 read-write value. The assigned value is used, along with the value of
16 aAggPortPartnerAdminSystemPriority, aAggPortPartnerAdminSystemID,
17 aAggPortPartnerAdminPort, and aAggPortPartnerAdminPortPriority, in order to achieve
18 manually configured aggregation.
19
20

21 **7.3.2.1.11 aAggPortPartnerOperKey**

22 ATTRIBUTE

23 APPROPRIATE SYNTAX:

24 INTEGER

25 BEHAVIOUR DEFINED AS:

26 The current operational value of the Key for the protocol Partner. The value of this attribute
27 may contain the manually configured value carried in aAggPortPartnerAdminKey if there is
28 no protocol Partner. This is a 16-bit read-only value.
29
30
31

32 **7.3.2.1.12 aAggPortSelectedAggID**

33 ATTRIBUTE

34 APPROPRIATE SYNTAX:

35 INTEGER

36 BEHAVIOUR DEFINED AS:

37 The identifier value of the Aggregator that this Aggregation Port has currently selected. Zero
38 indicates that the Aggregation Port has not selected an Aggregator, either because it is in the
39 process of detaching from an Aggregator or because there is no suitable Aggregator available
40 for it to select. This value is read-only.
41
42
43

44 **7.3.2.1.13 aAggPortAttachedAggID**

45 ATTRIBUTE

46 APPROPRIATE SYNTAX:

47 INTEGER

48 BEHAVIOUR DEFINED AS:

49 The identifier value of the Aggregator to which this Aggregation Port is currently attached.
50 Zero indicates that the Aggregation Port is not currently attached to an Aggregator. This
51 value is read-only.
52
53
54

7.3.2.1.14 aAggPortActorPort

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The Port Number locally assigned to the Aggregation Port. The Port Number is communicated in LACPDUs as the Actor_Port. This value is read-only.

7.3.2.1.15 aAggPortActorPortPriority

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The priority value assigned to this Aggregation Port. This 16-bit value is read-write.

NOTE—In the case of DRNI (Clause 9), the two least significant bits of the priority for each Aggregation Port in a Distributed Relay's Aggregator Port will be ignored because these bits are used to encode the Portal System Number [item e) in 9.3.4].

7.3.2.1.16 aAggPortPartnerAdminPort

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The current administrative value of the Port Number for the protocol Partner. This is a 16-bit read-write value. The assigned value is used, along with the value of aAggPortPartnerAdminSystemPriority, aAggPortPartnerAdminSystemID, aAggPortPartnerAdminKey, and aAggPortPartnerAdminPortPriority, in order to achieve manually configured aggregation.

7.3.2.1.17 aAggPortPartnerOperPort

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The operational Port Number assigned to this Aggregation Port by the Aggregation Port's protocol Partner. The value of this attribute may contain the manually configured value carried in aAggPortPartnerAdminPort if there is no protocol Partner. This 16-bit value is read-only.

7.3.2.1.18 aAggPortPartnerAdminPortPriority

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

1 The current administrative value of the Port Priority for the protocol Partner. This is a 16-bit
2 read-write value. The assigned value is used, along with the value of
3 aAggPortPartnerAdminSystemPriority, aAggPortPartnerAdminSystemID,
4 aAggPortPartnerAdminKey, and aAggPortPartnerAdminPort, in order to achieve manually
5 configured aggregation.
6

7 7.3.2.1.19 aAggPortPartnerOperPortPriority

8 ATTRIBUTE

9 APPROPRIATE SYNTAX:

10 INTEGER

11 BEHAVIOUR DEFINED AS:

12 The priority value assigned to this Aggregation Port by the Partner. The value of this attribute
13 may contain the manually configured value carried in aAggPortPartnerAdminPortPriority if
14 there is no protocol Partner. This 16-bit value is read-only.
15
16

17 7.3.2.1.20 aAggPortActorAdminState

18 ATTRIBUTE

19 APPROPRIATE SYNTAX:

20 BIT STRING [SIZE (1..8)]

21 BEHAVIOUR DEFINED AS:

22 A string of 8 bits, corresponding to the administrative values of Actor_State (6.4.2) as
23 transmitted by the Actor in LACPDU. The first bit corresponds to bit 0 of Actor_State
24 (LACP_Activity), the second bit corresponds to bit 1 (LACP_Timeout), the third bit
25 corresponds to bit 2 (Aggregation), the fourth bit corresponds to bit 3 (Synchronization), the
26 fifth bit corresponds to bit 4 (Collecting), the sixth bit corresponds to bit 5 (Distributing), the
27 seventh bit corresponds to bit 6 (Defaulted), and the eighth bit corresponds to bit 7 (Expired).
28 These values allow administrative control over the values of LACP_Activity,
29 LACP_Timeout, and Aggregation. This attribute value is read-write.
30
31
32

33 7.3.2.1.21 aAggPortActorOperState

34 ATTRIBUTE

35 APPROPRIATE SYNTAX:

36 BIT STRING [SIZE (1..8)]

37 BEHAVIOUR DEFINED AS:

38 A string of 8 bits, corresponding to the current operational values of Actor_State (6.4.2) as
39 transmitted by the Actor in LACPDU. The bit allocations are as defined in 7.3.2.1.20. This
40 attribute value is read-only.
41
42

43 7.3.2.1.22 aAggPortPartnerAdminState

44 ATTRIBUTE

45 APPROPRIATE SYNTAX:

46 BIT STRING [SIZE (1..8)]

47 BEHAVIOUR DEFINED AS:

48 A string of 8 bits, corresponding to the current administrative value of Actor_State for the
49 protocol Partner. The bit allocations are as defined in 7.3.2.1.20. This attribute value is read-
50 write. The assigned value is used in order to achieve manually configured aggregation.
51
52
53
54

7.3.2.1.23 aAggPortPartnerOperState

ATTRIBUTE

APPROPRIATE SYNTAX:

BIT STRING [SIZE (1..8)]

BEHAVIOUR DEFINED AS:

A string of 8 bits, corresponding to the current values of Actor_State in the most recently received LACPDU transmitted by the protocol Partner. The bit allocations are as defined in 7.3.2.1.20. In the absence of an active protocol Partner, this value may reflect the manually configured value aAggPortPartnerAdminState. This attribute value is read-only.

7.3.2.1.24 aAggPortAggregateOrIndividual

ATTRIBUTE

APPROPRIATE SYNTAX:

BOOLEAN

BEHAVIOUR DEFINED AS:

A read-only Boolean value indicating whether the Aggregation Port is able to Aggregate (“TRUE”) or is only able to operate as an Individual link (“FALSE”).

7.3.2.1.25 aAggPortOperConversationPasses

ATTRIBUTE

APPROPRIATE SYNTAX

BIT STRING [SIZE (4096)]

BEHAVIOR DEFINED AS

A read-only current operational vector of Boolean values, with one value for each possible Port Conversation ID. A 1 indicates that the Port Conversation ID is distributed through this Aggregation Port, and a 0 indicates that it cannot. aAggPortOperConversationPasses is referencing the current value of Port_Oper_Conversation_Mask (6.6.2.2).

7.3.2.1.26 aAggPortOperConversationCollected

ATTRIBUTE

APPROPRIATE SYNTAX

BIT STRING [SIZE (4096)]

BEHAVIOR DEFINED AS

A read-only current operational vector of Boolean values, with one value for each possible Port Conversation ID. A 1 indicates that the Port Conversation ID is collected through this Aggregation Port, and a 0 indicates that it cannot. aAggPortOperConversationPasses is referencing the current value of Collection_Conversation_Mask (6.6.1.1.2).

7.3.2.1.27 aAggPortLinkNumberID

ATTRIBUTE

APPROPRIATE SYNTAX

INTEGER, 0 to 65535

BEHAVIOUR DEFINED AS:

The Link Number ID value configured for this Aggregation Port by the System’s administrator. When the Link Number ID value matches one of the non zero values in the

1 selection prioritized lists in aAggConversationAdminLink[] (7.3.1.1.35), then this
2 Aggregation Port must be configured to have an aAggPortActorAdminKey value that
3 matches the aAggActorAdminKey of the Aggregator used by the LAG of the links specified
4 in aAggConversationAdminLink[]. Its default value is set to aAggPortActorPort (7.3.2.1.14).
5

6 NOTE—In the case of DRNI, the match of the aAggActorAdminKey to aAggPortActorAdminKey values excludes the
7 first two bits identifying the individual Portal System in the Portal. If the network administrator fails to configure the
8 proper values for the aAggPortActorAdminKey variables in all of the Aggregators Ports attached to a Portal, the
9 Distributed Relay Control Protocol (DRCP, 9.4) and the variable Port_Oper_Conversation_Mask (6.6.2.2) prevent
10 looping and/or duplicate delivery, if necessary, by discarding frames belonging to misconfigured Conversations.
11

12 **7.3.2.1.28 aAggPortPartnerAdminLinkNumberID**

13 ATTRIBUTE

14 APPROPRIATE SYNTAX

15 INTEGER, 0 to 65535

16 BEHAVIOUR DEFINED AS:

17 The value for the Link Number ID of the Partner System for this Aggregation Port, assigned
18 by administrator or System policy for use when the Partner's information is unknown. Its
19 default value is set to 0.
20
21

22 **7.3.2.1.29 aAggPortWTRTime**

23 ATTRIBUTE

24 APPROPRIATE SYNTAX

25 INTEGER

26 BEHAVIOUR DEFINED AS:

27 The wait-to-restore (WTR) period accompanying selection rules set by
28 aAggConversationAdminLink[] in a command issued by a network administrator. It may be
29 configured in steps of 1 min between 5 min and 12 min, while two additional special values
30 are also used. The value 0 indicates revertive and is the default value. The value 100 indicates
31 non-revertive mode of operation and the WTR_timer will be kept to the value 100.
32
33
34

35 **7.3.2.2 Aggregation Port Extension Attributes**

36 **7.3.2.2.1 aAggPortProtocolDA**

37 ATTRIBUTE

38 APPROPRIATE SYNTAX

39 MACAddress

40 BEHAVIOR DEFINED AS

41 A 6-octet read-write MACAddress value specifying the destination address to be used when
42 sending Link Aggregation Control and Marker PDUs on this Aggregation Port,
43 corresponding to the value of Protocol_DA in 6.2.8.1.2, 6.2.10.1.3 and 6.5.4.2.1. The default
44 value shall be the IEEE 802.3 Slow_Protocols_Multicast address.
45
46
47

48 **7.3.3 Aggregation Port Statistics managed object class**

49 This subclause formally defines the behaviours for the oAggPortStats managed object class attributes.
50
51
52
53
54

7.3.3.1 Aggregation Port Statistics attributes

7.3.3.1.1 aAggPortStatsID

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

This read-only attribute identifies an Aggregation Port Statistics object instance among the subordinate managed objects of the containing object. The value allocated to this attribute shall be the same as the containing oAggregationPort managed object.

7.3.3.1.2 aAggPortStatsLACPDUsRx

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

The number of valid LACPDUs received on this Aggregation Port. This value is read-only.

7.3.3.1.3 aAggPortStatsMarkerPDUsRx

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

The number of valid Marker PDUs received on this Aggregation Port. This value is read-only.

7.3.3.1.4 aAggPortStatsMarkerResponsePDUsRx

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

The number of valid Marker Response PDUs received on this Aggregation Port. This value is read-only.

7.3.3.1.5 aAggPortStatsUnknownRx

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 50 counts per second.

BEHAVIOUR DEFINED AS:

The number of frames received that either

—Carry the Slow Protocols Ethernet Type value (IEEE Std 802.3 Annex 57A.4), but contain an unknown PDU, or

—Are addressed to the Slow Protocols group MAC Address (IEEE Std 802.3 Annex 57A.3), but do not carry the Slow Protocols Ethernet Type. This value is read-only.

7.3.3.1.6 aAggPortStatsIllegalRx

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 50 counts per second.

BEHAVIOUR DEFINED AS:

The number of frames received that carry the Slow Protocols Ethernet Type value (IEEE Std 802.3 Annex 57A.4), but contain a badly formed PDU or an illegal value of Protocol Subtype (IEEE Std 802.3 Annex 57A.3). This value is read-only.

7.3.3.1.7 aAggPortStatsLACPDUsTx

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

The number of LACPDUs transmitted on this Aggregation Port. This value is read-only.

7.3.3.1.8 aAggPortStatsMarkerPDUsTx

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

The number of Marker PDUs transmitted on this Aggregation Port. This value is read-only.

7.3.3.1.9 aAggPortStatsMarkerResponsePDUsTx

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

The number of Marker Response PDUs transmitted on this Aggregation Port. This value is read-only.

7.3.4 Aggregation Port Debug Information managed object class

This subclause formally defines the behaviours for the oAggPortDebugInformation managed object class attributes.

7.3.4.1 Aggregation Port Debug Information attributes

7.3.4.1.1 aAggPortDebugInformationID

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

1 This read-only attribute identifies an LACP Debug Information object instance among the
2 subordinate managed objects of the containing object. The value allocated to this attribute
3 shall be the same as the containing oAggregationPort managed object.
4

5 **7.3.4.1.2 aAggPortDebugRxState**

6 ATTRIBUTE

7 APPROPRIATE SYNTAX:

8 An ENUMERATED VALUE that has one of the following entries:

9 current

10 expired

11 defaulted

12 initialize

13 lacpDisabled

14 portDisabled

15 BEHAVIOUR DEFINED AS:

16 This attribute holds the value “current” if the Receive state machine for the Aggregation Port
17 is in the CURRENT state, “expired” if the Receive state machine is in the EXPIRED state,
18 “defaulted” if the Receive state machine is in the DEFAULTED state, “initialize” if the
19 Receive state machine is in the INITIALIZE state, “lacpDisabled” if the Receive state
20 machine is in the LACP_DISABLED state, or “portDisabled” if the Receive state machine is
21 in the PORT_DISABLED state. This value is read-only.
22
23

24 **7.3.4.1.3 aAggPortDebugLastRxTime**

25 ATTRIBUTE

26 APPROPRIATE SYNTAX:

27 INTEGER

28 BEHAVIOUR DEFINED AS:

29 The time at which the last LACPDU was received by this Aggregation Port, in terms of
30 centiseconds since the system was last reset. The ifLastChange object in the Interfaces MIB
31 defined in IETF RFC 2863 is a suitable object for supplying a value for
32 aAggPortDebugLastRxTime. This value is read-only.
33
34
35

36 NOTE—aAggPortDebugLastRxTime was defined in terms of the aTimeSinceSystemReset variable of IEEE Std 802.3
37 Annex F.2.1 in earlier versions of this standard. aTimeSinceSystemReset and ifLastChange have the same meaning.
38

39 **7.3.4.1.4 aAggPortDebugMuxState**

40 ATTRIBUTE

41 APPROPRIATE SYNTAX:

42 An ENUMERATED VALUE that has one of the following entries:

43 detached

44 waiting

45 attached

46 collecting

47 distributing

48 collecting_distributing

49 BEHAVIOUR DEFINED AS:

50 This attribute holds the value “detached” if the Mux state machine (6.4.15) for the
51 Aggregation Port is in the DETACHED state, “waiting” if the Mux state machine for the
52
53
54

1 Aggregation Port is in the WAITING state, “attached” if the Mux state machine for the
2 Aggregation Port is in the ATTACHED state, “collecting” if the Mux state machine for the
3 Aggregation Port is in the COLLECTING state, “distributing” if the Mux state machine for the
4 Aggregation Port is in the DISTRIBUTING state, and “collecting_distributing” if the
5 Mux state machine for the Aggregation Port is in the COLLECTING_DISTRIBUTING state.
6 This value is read-only.
7

8 **7.3.4.1.5 aAggPortDebugMuxReason**

9
10 ATTRIBUTE

11 APPROPRIATE SYNTAX:

12 A PrintableString, 255 characters max.

13 BEHAVIOUR DEFINED AS:

14 A human-readable text string indicating the reason for the most recent change of Mux
15 machine state. This value is read-only.
16

17 **7.3.4.1.6 aAggPortDebugActorChurnState**

18
19 ATTRIBUTE

20 APPROPRIATE SYNTAX:

21 An ENUMERATED VALUE that has one of the following entries:

22 noChurn

23 churn

24
25 BEHAVIOUR DEFINED AS:

26 The state of the Actor Churn Detection machine (6.4.17) for the Aggregation Port. A value of
27 “noChurn” indicates that the state machine is in either the NO_ACTOR_CHURN or the
28 ACTOR_CHURN_MONITOR state, and “churn” indicates that the state machine is in the
29 ACTOR_CHURN state. This value is read-only.
30

31 **7.3.4.1.7 aAggPortDebugPartnerChurnState**

32
33 ATTRIBUTE

34 APPROPRIATE SYNTAX:

35 An ENUMERATED VALUE that has one of the following entries:

36 noChurn

37 churn

38
39 BEHAVIOUR DEFINED AS:

40 The state of the Partner Churn Detection machine (6.4.17) for the Aggregation Port. A value
41 of “noChurn” indicates that the state machine is in either the NO_PARTNER_CHURN or the
42 PARTNER_CHURN_MONITOR state, and “churn” indicates that the state machine is in the
43 PARTNER_CHURN state. This value is read-only.
44

45 **7.3.4.1.8 aAggPortDebugActorChurnCount**

46
47 ATTRIBUTE

48 APPROPRIATE SYNTAX:

49 Generalized counter. This counter has a maximum increment rate of 5 counts per second.

50 BEHAVIOUR DEFINED AS:

51 Count of the number of times the Actor Churn state machine has entered the
52 ACTOR_CHURN state. This value is read-only.
53
54

7.3.4.1.9 aAggPortDebugPartnerChurnCount

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

Count of the number of times the Partner Churn state machine has entered the PARTNER_CHURN state. This value is read-only.

7.3.4.1.10 aAggPortDebugActorSyncTransitionCount

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

Count of the number of times the Actor's Mux state machine (6.4.15) has entered the IN_SYNC state. This value is read-only.

7.3.4.1.11 aAggPortDebugPartnerSyncTransitionCount

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

Count of the number of times the Partner's Mux state machine (6.4.15) has entered the IN_SYNC state. This value is read-only.

7.3.4.1.12 aAggPortDebugActorChangeCount

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

Count of the number of times the Actor's perception of the LAG ID for this Aggregation Port has changed. This value is read-only.

7.3.4.1.13 aAggPortDebugPartnerChangeCount

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized counter. This counter has a maximum increment rate of 5 counts per second.

BEHAVIOUR DEFINED AS:

Count of the number of times the Partner's perception of the LAG ID (6.3.6.1) for this Aggregation Port has changed. This value is read-only.

7.3.4.1.14 aAggPortDebugActorCDSChurnState

ATTRIBUTE

APPROPRIATE SYNTAX:

An ENUMERATED VALUE that has one of the following entries:

noChurn

1 churn

2 BEHAVIOUR DEFINED AS:

3 This managed object is applicable only when Conversation-sensitive frame collection and
4 distribution as specified in 6.6 is supported. The state of the Actor CDS Churn Detection
5 machine (6.6.2.7) for the Aggregation Port. A value of “noChurn” indicates that the state
6 machine is in either the NO_ACTOR_CDS_CHURN or the
7 ACTOR_CHURN_CDS_MONITOR state, and “churn” indicates that the state machine is in
8 the ACTOR_CDS_CHURN state. This value is read-only.
9

10 **7.3.4.1.15 aAggPortDebugPartnerCDSChurnState**

11 ATTRIBUTE

12 APPROPRIATE SYNTAX:

13 An ENUMERATED VALUE that has one of the following entries:

14 noChurn

15 churn

16 BEHAVIOUR DEFINED AS:

17 This managed object is applicable only when Conversation-sensitive frame collection and
18 distribution as specified in 6.6 is supported. The state of the Partner CDS Churn Detection
19 machine (6.6.2.7) for the Aggregation Port. A value of “noChurn” indicates that the state
20 machine is in either the NO_PARTNER_CDS_CHURN or the
21 PARTNER_CDS_CHURN_MONITOR state, and “churn” indicates that the state machine is
22 in the PARTNER_CDSCHURN state. This value is read-only.
23
24

25 **7.3.4.1.16 aAggPortDebugActorCDSChurnCount**

26 ATTRIBUTE

27 APPROPRIATE SYNTAX:

28 Generalized counter. This counter has a maximum increment rate of 5 counts per second.

29 BEHAVIOUR DEFINED AS:

30 This managed object is applicable only when Conversation-sensitive frame collection and
31 distribution as specified in 6.6 is supported. Count of the number of times the Actor CDS
32 Churn state machine has entered the ACTOR_CDS_CHURN state. This value is read-only.
33
34
35

36 **7.3.4.1.17 aAggPortDebugPartnerCDSChurnCount**

37 ATTRIBUTE

38 APPROPRIATE SYNTAX:

39 Generalized counter. This counter has a maximum increment rate of 5 counts per second.

40 BEHAVIOUR DEFINED AS:

41 This managed object is applicable only when Conversation-sensitive frame collection and
42 distribution as specified in 6.6 is supported. Count of the number of times the Partner CDS
43 Churn state machine has entered the PARTNER_CDS_CHURN state. This value is read-
44 only.
45
46
47

48 **7.4 Management for Distributed Resilient Network Interconnect**

49 **7.4.1 Distributed Relay Managed Object Class**

50 This subclause formally defines the behaviours for the oDistributedRelay managed object class attributes.
51
52
53
54

7.4.1.1 Distributed Relay Attributes

7.4.1.1.1 aDrniID

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The unique identifier allocated to this Distributed Relay by the local System. This attribute identifies a Distributed Relay instance among the subordinate managed objects of the containing object. This value is read-only.

NOTE—The aDrniID is represented in the SMIV2 MIB as an ifIndex—see D.5.

7.4.1.1.2 aDrniDescription

ATTRIBUTE

APPROPRIATE SYNTAX:

A PrintableString, 255 characters max.

BEHAVIOUR DEFINED AS:

A human-readable text string containing information about the Distributed Relay. This string is read-only. The contents are vendor specific.

7.4.1.1.3 aDrniName

ATTRIBUTE

APPROPRIATE SYNTAX:

A PrintableString, 255 characters max.

BEHAVIOUR DEFINED AS:

A human-readable text string containing a locally significant name for the Distributed Relay. This string is read-write.

7.4.1.1.4 aDrniPortalAddr

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE OF 6 OCTETS that match the syntax of a 48-bit MAC Address

BEHAVIOUR DEFINED AS:

A read-write identifier of a particular Portal. aDrniPortalAddr has to be unique among at least all of the potential Portal Systems to which a given Portal System might be attached via an IPL. Also used as the Actor's System ID (6.3.2) for the emulated system.

7.4.1.1.5 aDrniPortalPriority

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

A 2-octet read-write value indicating the priority value associated with the Portal's System ID. Also used as the Actor's System Priority (6.3.2) for the emulated system.

7.4.1.1.6 aDrniThreePortalSystem

ATTRIBUTE

APPROPRIATE SYNTAX:

BOOLEAN

BEHAVIOUR DEFINED AS:

A read-write Boolean value indicating whether this Portal System is part of a Portal consisting of three Portal Systems or not. Value 1 stands for a Portal of three Portal Systems, value 0 stands for a Portal of two or one Portal Systems. The default value is 0.

7.4.1.1.7 aDrniPortalSystemNumber

ATTRIBUTE

APPROPRIATE SYNTAX

A Portal System Number, which is an integer in the range 1 through 3 inclusive.

BEHAVIOR DEFINED AS

A read-write identifier of this particular Portal System within a Portal. It is the responsibility of the network administrator to ensure that these numbers are unique among the Portal Systems with the same aDrniPortalAddr (7.4.1.1.4).

7.4.1.1.8 aDrniIntraPortalLinkList

ATTRIBUTE

APPROPRIATE SYNTAX

A SEQUENCE OF INTEGERS that match the syntax of an Interface Identifier.

BEHAVIOR DEFINED AS

Read-write list of the Interface Identifiers of the Ports to the Intra-Portal Links assigned to this Distributed Relay. Each Interface Identifier, a Port ID (6.3.4), has the two least significant bits of its Port Priority (7.3.2.1.15) configured to match the Portal System Number of the attached Portal System. The number of IPLs in the list depends on the Portal topology. For a Portal of three Portal Systems two or three IPLs can be used, for a Portal of two Portal Systems a single IPL is required and for a single Portal System no IPL is required.

7.4.1.1.9 aDrniAggregator

ATTRIBUTE

APPROPRIATE SYNTAX

An INTEGER that matches the syntax of an Interface Identifier.

BEHAVIOR DEFINED AS

Read-write Interface Identifier of the Aggregator Port assigned to this Distributed Relay.

7.4.1.1.10 aDrniConvAdminGateway[]

ATTRIBUTE

APPROPRIATE SYNTAX

An array of SEQUENCE OF INTEGERS that match the syntax of Portal System Number.

BEHAVIOR DEFINED AS

There are 4096 aDrniConvAdminGateway[] variables, aDrniConvAdminGateway[0] through aDrniConvAdminGateway[4095], indexed by Gateway Conversation ID. Each contains administrative values of the Gateway selection priority list for the Distributed Relay for the referenced Gateway Conversation ID. This selection priority list, a sequence of

1 integers for each Gateway Conversation ID, is a list of Portal System Numbers in the order of
2 preference, highest to lowest, for the corresponding preferred Portal System's Gateway to
3 carry that Conversation.
4

5 NOTE—To the extent that the network administrator fails to configure the same values for the
6 aDrniConvAdminGateway[] variables in all of the DR Functions of a Portal, frames can be misdirected. The Distributed
7 Relay Control Protocol (DRCP, 9.4) detects such misconfiguration.
8

9 **7.4.1.1.11 aDrniNeighborAdminConvGatewayListDigest**

10 ATTRIBUTE

11 APPROPRIATE SYNTAX

12 A SEQUENCE OF OCTETS consisting of a 16-octet MD5 Digest.

13 BEHAVIOR DEFINED AS

14 The value for the digest of the prioritized Gateway Conversation ID-to-Gateway assignments
15 of the Neighbor Portal System, assigned by administrator or System policy for use when the
16 Neighbor Portal System's information is unknown. Its default value is set to NULL.
17
18

19 **7.4.1.1.12 aDrniNeighborAdminConvPortListDigest**

20 ATTRIBUTE

21 APPROPRIATE SYNTAX

22 A SEQUENCE OF OCTETS consisting of a 16-octet MD5 Digest.

23 BEHAVIOR DEFINED AS

24 The value for the digest of the prioritized Port Conversation ID-to-Aggregation Port
25 assignments of the Neighbor Portal System, assigned by administrator or System policy for
26 use when the Neighbor Portal System's information is unknown. Its default value is set to
27 NULL.
28
29

30 **7.4.1.1.13 aDrniGatewayAlgorithm**

31 ATTRIBUTE

32 APPROPRIATE SYNTAX

33 A SEQUENCE OF OCTETS consisting of an OUI or CID and one following octet.

34 BEHAVIOR DEFINED AS

35 This object identifies the algorithm used by the DR Function to assign frames to a Gateway
36 Conversation ID. Table 9-7 provides the IEEE 802.1 OUI (00-80-C2) Gateway Algorithm
37 encodings.
38
39
40

41 **7.4.1.1.14 aDrniNeighborAdminGatewayAlgorithm**

42 ATTRIBUTE

43 APPROPRIATE SYNTAX

44 A SEQUENCE OF OCTETS consisting of a three-octet OUI or CID and one following octet.

45 BEHAVIOR DEFINED AS

46 This object identifies the value for the Gateway algorithm of the Neighbor Portal System,
47 assigned by administrator or System policy for use when the Neighbor Portal System's
48 information is unknown. Table 9-7 provides the IEEE 802.1 OUI (00-80-C2) Gateway
49 Algorithm encodings. Its default value is set to NULL.
50
51
52
53
54

7.4.1.1.15 aDrniNeighborAdminPortAlgorithm

ATTRIBUTE

APPROPRIATE SYNTAX

A SEQUENCE OF OCTETS consisting of a three-octet OUI or CID and one following octet.

BEHAVIOR DEFINED AS

This object identifies the value for the Port Algorithm of the Neighbor Portal System, assigned by administrator or System policy for use when the Neighbor Portal System's information is unknown. Table 6-4 provides the IEEE 802.1 OUI (00-80-C2) Port Algorithm encodings. Its default value is set to NULL.

7.4.1.1.16 aDrniNeighborAdminDRCPState

ATTRIBUTE

APPROPRIATE SYNTAX:

BIT STRING [SIZE (1..8)]

BEHAVIOUR DEFINED AS:

A string of 8 bits, corresponding to the administrative values of DRCP_State [item s] in 9.4.3.2] as transmitted by this Portal System in DRCPDUs. The first bit corresponds to bit 0 of DRCP_State (Home_Gateway), the second bit corresponds to bit 1 (Neighbor_Gateway), the third bit corresponds to bit 2 (Other_Gateway), the fourth bit corresponds to bit 3 (IPP_Activity), the fifth bit corresponds to bit 4 (DRCP_Timeout), the sixth bit corresponds to bit 5 (Gateway_Sync), the seventh bit corresponds to bit 6 (Port_Sync), and the eighth bit corresponds to bit 7 (Expired). These values allow administrative control over the values of Home_Gateway, Neighbor_Gateway, Other_Gateway, IPP_Activity, and DRCP_Timeout. Their values are by default set to FALSE. This attribute value is read-write.

7.4.1.1.17 aDrniEncapsulationMethod

ATTRIBUTE

APPROPRIATE SYNTAX

A SEQUENCE OF OCTETS consisting of an OUI or CID and one following octet.

BEHAVIOR DEFINED AS

This managed object is applicable only when Network / IPL sharing by time (9.3.2.1) or Network / IPL sharing by tag (9.3.2.2) or Network / IPL sharing by encapsulation (9.3.2.3) is supported. The object identifies the value representing the encapsulation method that is used to transport IPL frames to the Neighbor Portal System when the IPL and network link are sharing the same physical link. It consists of the three-octet OUI or CID identifying the organization which is responsible for this encapsulation and one following octet used to identify the encapsulation method defined by that organization. Table 9-11 provides the IEEE 802.1 OUI (00-80-C2) encapsulation method encodings. A Default value of 0x00-80-C2-00 indicates that the IPL is using a separate physical or Aggregation link. A value of 1 indicates that Network / IPL sharing by time (9.3.2.1) is used. A value of 2 indicates that the encapsulation method used is the same as the one used by network frames and that Network / IPL sharing by tag (9.3.2.2) is used.

7.4.1.1.18 aDrniPLEncapMap

ATTRIBUTE

APPROPRIATE SYNTAX

A SEQUENCE OF INTEGERS, indexed by Gateway Conversation ID.

BEHAVIOR DEFINED AS

1 This managed object is applicable only when Network / IPL sharing by tag (9.3.2.2) or
2 Network / IPL sharing by encapsulation (9.3.2.3) is supported. Each entry represents the
3 value of the identifier used for an IPL frame associated with that Gateway Conversation ID
4 for the encapsulation method specified in 7.4.1.1.17.
5

6 **7.4.1.1.19 aDrniNetEncapMap**

7 ATTRIBUTE

8 APPROPRIATE SYNTAX

9 A SEQUENCE OF INTEGERS, indexed by Gateway Conversation ID.

10 BEHAVIOR DEFINED AS

11 This managed object is applicable only when Network / IPL sharing by tag (9.3.2.2) is
12 supported. Each entry represents the translated value of the identifier used for a network
13 frame associated with that Gateway Conversation ID when the method specified in 7.4.1.1.17
14 is the Network / IPL sharing by tag method specified in 9.3.2.2 and the network frames need
15 to share the tag space used by IPL frames.
16
17

18 **7.4.1.1.20 aDrniDRPortConversationPasses**

19 ATTRIBUTE

20 APPROPRIATE SYNTAX

21 BIT STRING [SIZE (4096)]

22 BEHAVIOR DEFINED AS

23 A read-only current operational vector of Boolean values, with one value for each possible
24 Port Conversation ID. A 1 indicates that the Port Conversation ID is allowed to be distributed
25 through this DR Function's Aggregator, and a 0 indicates that it cannot.
26 aDrniDRPortConversationPasses is referencing the current value of
27 Drni_Portal_System_Port_Conversation (9.3.4.2).
28
29
30

31 **7.4.1.1.21 aDrniDRGatewayConversationPasses**

32 ATTRIBUTE

33 APPROPRIATE SYNTAX

34 BIT STRING [SIZE (4096)]

35 BEHAVIOR DEFINED AS

36 A read-only current operational vector of Boolean values, with one value for each possible
37 Gateway Conversation ID. A 1 indicates that the Gateway Conversation ID is allowed to pass
38 through this DR Function's Gateway, and a 0 indicates that it cannot.
39 aDrniDRGatewayConversationPasses is referencing the current value of
40 Drni_Portal_System_Gateway_Conversation (9.3.4.2).
41
42
43

44 **7.4.1.1.22 aDrniPSI**

45 ATTRIBUTE

46 APPROPRIATE SYNTAX:

47 BOOLEAN

48 BEHAVIOUR DEFINED AS:

49 A read-only Boolean value providing the value of PSI, which indicates whether this Portal
50 System is isolated from the other Portal Systems within the same Portal ("TRUE") or not
51 ("FALSE").
52
53
54

7.4.1.1.23 aDrniPortConversationControl

ATTRIBUTE

APPROPRIATE SYNTAX:

BOOLEAN

BEHAVIOUR DEFINED AS:

A read-write Boolean value that controls the operation of the updateDRFHomeState (9.4.11). When set to “TRUE” the Home Gateway Vector is set equal to Drni_Portal_System_Port_Conversation. Setting this object to “TRUE” is only possible when the Gateway algorithm and the Port algorithm use the same distributions methods. The default is “FALSE”, indicating that the Home Gateway Vector is controlled by the network control protocol.

7.4.1.1.24 aDrniIntraPortalPortProtocolDA

ATTRIBUTE

APPROPRIATE SYNTAX

A SEQUENCE OF 6 OCTETS that match the syntax of a 48-bit MAC Address

BEHAVIOR DEFINED AS

A 6-octet read-write MAC Address value specifying the destination address to be used when sending DRCPDUs, corresponding to the value of DRCP_Protocol_DA in 9.4.4.1.3. Its values is one of the addresses selected from Table 9-6 and its default shall be the IEEE 802.1 Nearest non-TPMR Bridge group address (01-80-C2-00-00-03).

7.4.2 IPP Managed Objects Class

This subclause formally defines the behaviours for the oDistributedRelayIPP managed object class attributes.

7.4.2.1 IPP Attributes

7.4.2.1.1 aIPPID

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The unique identifier allocated to this IPP by the local Portal System. This attribute identifies an IPP instance among the subordinate managed objects of the containing object. This value is read-only.

NOTE—The aIPPID is represented in the SMiv2 MIB as an ifIndex—see D.5.

7.4.2.1.2 aIPPPortConversationPasses

ATTRIBUTE

APPROPRIATE SYNTAX

BIT STRING [SIZE (4096)]

BEHAVIOR DEFINED AS

A read-only current operational vector of Boolean values, with one value for each possible Port Conversation ID. A 1 indicates that the Port Conversation ID is allowed to be transmitted through this IPP, and a 0 indicates that it cannot. aIPPPortConversationPasses is referencing the current value of Ipp_Port_Conversation_Passes (9.3.4.3).

7.4.2.1.3 aIPPGatewayConversationDirection

ATTRIBUTE

APPROPRIATE SYNTAX

BIT STRING [SIZE (4096)]

BEHAVIOR DEFINED AS

A read-only current operational vector of Boolean values, with one value for each possible Gateway Conversation ID. A 1 indicates that the Gateway Conversation ID is assigned to Gateways reachable through this IPP, and a 0 indicates that the Gateway for the indexed Gateway Conversation ID is not reachable through this IPP. aIPPGatewayConversationDirection is referencing the current value of Ipp_Gateway_Conversation_Direction (9.3.4.3).

7.4.2.1.4 aIPPAdminState

ATTRIBUTE

APPROPRIATE SYNTAX:

An ENUMERATED VALUE that has one of the following entries:

up
down

BEHAVIOUR DEFINED AS:

This read-write value defines the administrative state of the IPP. A value of “up” indicates that the operational state of the IPP (aIPPOperState) is permitted to be either “up” or “down”. A value of “down” forces the operational state of the IPP to be “down”. A GET operation returns the current administrative state. A SET operation changes the administrative state to a new value.

7.4.2.1.5 aIPPOperState

ATTRIBUTE

APPROPRIATE SYNTAX:

An ENUMERATED VALUE that has one of the following entries:

up
down

BEHAVIOUR DEFINED AS:

This read-only value defines the operational state of the IPP. The operational state is “up” if the IPL is operational, and if the value of aIPPAdminState for the IPP is also “up.” If the IPL is not operational, or if the administrative state of the IPP (aIPPAdminState) is “down”, then the operational state is “down.” An operational state of “up” indicates that the IPP is available for use by the DR Function; a value of “down” indicates that the IPP is not available for use by the DR Function.

7.4.2.1.6 aIPPTimeOfLastOperChange

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The time at which the interface entered its current operational state, in terms of centiseconds since the system was last reset. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this object contains a value of

1 zero. The ifLastChange object in the Interfaces MIB defined in IETF RFC 2863 is a suitable
2 object for supplying a value for aIPPTimeOfLastOperChange. This value is read-only.
3
4

5 **7.4.3 IPP Statistics managed object class**

6
7 This subclause formally defines the behaviours for the oIPPStats managed object class attributes.
8

9 **7.4.3.1 IPP Statistics attributes**

10 **7.4.3.1.1 aIPPStatsID**

11
12
13 ATTRIBUTE

14 APPROPRIATE SYNTAX:

15 INTEGER

16 BEHAVIOUR DEFINED AS:

17 This read-only attribute identifies an IPP Statistics object instance among the subordinate
18 managed objects of the containing object. The value allocated to this attribute shall be the
19 same as the containing oDistributedRelayIPP managed object.
20

21 **7.4.3.1.2 aIPPStatsDRCPDUsRx**

22
23 ATTRIBUTE

24 APPROPRIATE SYNTAX:

25 Generalized counter. This counter has an expected increment rate of 5 counts per second.

26 BEHAVIOUR DEFINED AS:

27 The number of valid DRCPDUs received on this IPP. This value is read-only.
28
29

30 **7.4.3.1.3 aIPPStatsIllegalRx**

31
32 ATTRIBUTE

33 APPROPRIATE SYNTAX:

34 Generalized counter. This counter has an expected increment rate of 50 counts per second.

35 BEHAVIOUR DEFINED AS:

36 The number of frames received that carry the DRCP Ethernet Type value (9.4.2.4), but
37 contain a badly formed PDU. This value is read-only.
38
39

40 **7.4.3.1.4 aIPPStatsDRCPDUsTx**

41
42 ATTRIBUTE

43 APPROPRIATE SYNTAX:

44 Generalized counter. This counter has an expected increment rate of 5 counts per second.

45 BEHAVIOUR DEFINED AS:

46 The number of DRCPDUs transmitted on this IPP. This value is read-only.
47

48 **7.4.4 IPP Debug Information managed object class**

49
50 This subclause formally defines the behaviours for the oIPPDebugInformation managed object class
51 attributes.
52
53
54

7.4.4.1 IPP Debug Information attributes

7.4.4.1.1 aIPPDebugInformationID

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

This read-only attribute identifies an DRCP Debug Information object instance among the subordinate managed objects of the containing object. The value allocated to this attribute shall be the same as the containing oDistributedRelayIPP managed object.

7.4.4.1.2 aIPPDebugDRCPRxState

ATTRIBUTE

APPROPRIATE SYNTAX:

An ENUMERATED VALUE that has one of the following entries:

current

expired

defaulted

initialize

reportToManagement

BEHAVIOUR DEFINED AS:

This attribute holds the value “current” if the DRCPDU Receive state machine for the IPP is in the CURRENT state, “expired” if the DRCPDU Receive state machine is in the EXPIRED state, “defaulted” if the DRCPDU Receive state machine is in the DEFAULTED state, “initialize” if the DRCPDU Receive state machine is in the INITIALIZE state, or “reportToManagement” if the Receive state machine is in the REPORT_TO_MANAGEMENT state. This value is read-only.

7.4.4.1.3 aIPPDebugLastRxTime

ATTRIBUTE

APPROPRIATE SYNTAX:

INTEGER

BEHAVIOUR DEFINED AS:

The time at which the last DRCPDU was received by this IPP, in terms of centiseconds since the system was last reset. The ifLastChange object in the Interfaces MIB defined in IETF RFC 2863 is a suitable object for supplying a value for aIPPDebugLastRxTime. This value is read-only.

7.4.4.1.4 aIPPDebugDifferPortalReason

ATTRIBUTE

APPROPRIATE SYNTAX:

A PrintableString, 255 characters max.

BEHAVIOUR DEFINED AS:

A human-readable text string indicating the most recent set of variables that are responsible for setting the variable Differ_Portal or Differ_Conf_Portal (9.4.8) on this IPP to TRUE. This value is read-only.

8. Frame distribution and collection algorithms

One specific algorithm for frame distribution and collection is defined, in 8.2.

8.1 Conversation Identifiers

The number of different conversations, each of which has to be delivered in order, can be very large. Some activities, particularly ensuring frame ordering using Conversation-sensitive frame collection and distribution (6.6), require that functional elements in different Systems maintain and exchange information applying to particular conversations. The enumeration of large numbers of conversations (e.g., conversations identified by a 24-bit I-SID) is impractical for a protocol such as LACP.

Therefore, a Conversation Identifier (or Conversation ID) is defined as a value in the range 0 through 4095. By administrative means, every possible conversation is assigned to a single Conversation ID value for each supported Conversation ID type. More than one conversation can be assigned to a Conversation ID. It is not necessary that every Conversation ID value have any conversations assigned to it. In this standard, several types of Conversation ID are specified for different uses.

8.2 Per-service frame distribution

A System may implement Per-service frame distribution, as defined in this section.

8.2.1 Goals and objectives

Distributing frames to physical links by service ID, as defined in 8.2, provides the following:

- a) **Connectivity Fault Management congruity**—The only thing that Connectivity Fault Management (CFM, IEEE Std 802.1Q Clause 18) PDUs have in common with the data that they protect is, in general, the service ID(s) that they share. Per-service frame distribution ensures that the CFM PDUs traverse the same physical links as their data.
- b) **Bidirectional congruity**—Providing a means for the two ends of an Aggregation Group to use the same physical link in both directions for a given service ensures that a failed link or a link experiencing an excessive error rate will affect the fewest possible number of services and in general provide support for protocols that have strict symmetry requirements on their transmit and receive paths e.g. the IEEE Std 1588™-2008 Precision Time Protocol (PTP).
- c) **Ingress predictability**—Ingress metering is often applied on a per-service basis. Confining a service to a single physical link localizes the meter, facilitating this process.

NOTE—An Aggregation System or Portal running Per-service frame distribution can interoperate with an Aggregation System or Portal not running Per-service frame distribution. However, only the Link Aggregation goals in 6.1.1 can be met, not the goals listed in this section.

8.2.2 Overview

A device implementing Link Aggregation, may implement Per-service frame distribution. When enabled, Per-service frame distribution determines how Port Conversation IDs are derived from frames by the Aggregator and, in DRNI systems, the Distributed Relay. Whether Per-service frame distribution is enabled on either end of a particular Link Aggregation Group is an administrative choice.

When enabled, Per-service frame distribution distributes frames to physical links according to their Service Identifier (Service ID). The following tag types defined by IEEE Std 802.1Q-2011 Table 9–1 are supported by Per-service frame distribution:

- a) Customer VLAN Tag (VID field, Clause 9.6 of IEEE Std 802.1Q-2011);
- b) Service VLAN Tag or Backbone VLAN Tag (VID field, Clause 9.6 of IEEE Std 802.1Q-2011); and
- c) Backbone Service Instance Tag (I-SID field, Clause 9.7 of IEEE Std 802.1Q-2011).

The active function (Frame Distributor, Frame Collector, or DR Function) examines each frame presented to it. If the first item in the M_UNITDATA.request or M_UNITDATA.indication mac_service_data_unit parameter is the one tag from the above list for which active function is configured, encapsulated as appropriate to the medium according to IEEE Std 802.1Q-2011 Clause 9.4, then the 12-bit VID field or 24-bit I-SID field of the tag is used to determine the Service ID. A value of zero is used if the first item in the mac_service_data_unit is not the appropriate tag.

A System supporting Per-service frame distribution (6.6) shall support Conversation-sensitive frame collection and distribution, and shall support classification by Customer VLAN Tag for C-tagged interfaces and classification by Service VLAN tag for S-tagged and B-tagged interfaces. A System supporting Per-service frame distribution may support service classification by Backbone Service Instance Tag. A Backbone Edge Bridge supporting Per-service frame distribution on its Provider Instance Ports (PIPs, Clause 6.10 of IEEE Std 802.1Q-2011) or Customer Backbone Ports (CBPs, Clause 6.10 of IEEE Std 802.1Q-2011) shall support service classification by Backbone Service Instance Tag on those ports. A System supporting Per service frame distribution may support classification by any other methods that are identified in aAggPortAlgorithm (7.3.1.1.33).

8.2.3 Port Conversation Identifiers

Per-service frame distribution is defined in terms of Port Conversation IDs (8.1). For 12-bit tags, both VLAN IDs and Port Conversation IDs have 4095 values available, so Port Conversation ID can be synonymous with tag value, and no mapping is needed. However, when 24-bit I-SIDs are used as Service IDs, mapping to Port Conversation IDs is necessary. In this case both Aggregators at the two ends of the LAG have to be configured with the same Service ID to Port Conversation ID mapping entries (7.3.1.1.38). Conversation-sensitive LACPDUs exchange digests of the tables containing these entries (6.4.2.4.4) to verify that the mapping table configuration is coordinated.

9. Distributed Resilient Network Interconnect

This Clause describes the elements of Distributed Resilient Network Interconnect (DRNI) as follows:

- a) An overview of DRNI is given in 9.1.
- b) The Distributed Relay is introduced in 9.2.
- c) The operation of the Distributed Relay is described in detail in 9.3.
- d) The Distributed Relay Control Protocol is described in 9.4.

The models of operation in this clause provide a basis for specifying the externally observable behavior of the operation, and are not intended to place additional constraints on implementations; these can adopt any internal model of operation compatible with the externally observable behavior specified. Conformance of equipment to this standard is purely in respect of observable protocol.

9.1 Goals and Objectives

As described in Clause 6, Link Aggregation creates a Link Aggregation Group that is a collection of one or more physical links that appears, to higher layers, to be a single logical link. The Link Aggregation Group has two ends, each terminating in an Aggregation System.

DRNI expands the concept of Link Aggregation so that, at either one or both ends of a Link Aggregation Group, the single Aggregation System is replaced by a Portal, composed from one to three Portal Systems.

The Distributed Resilient Network Interconnect, as defined in this Clause, provides the following:

- a) **Link Aggregation**—DRNI provides benefits of Link Aggregation, as listed in 6.1.1, item a through item n.
- b) **Portals**—Connections between two cooperating sets of Systems (two Portals) are supported, in contrast to Link Aggregation (item o) in 6.1.1), so that connectivity between two networks can be maintained despite the failure of an entire System (and its connected links) belonging to a Portal.
- c) **Compatibility**—A multi-System Portal can connect to a single-System Portal or to an Aggregation System compliant with Clause 6 or with previous versions of this Standard.
- d) **Administrative isolation**—A DRNI Link Aggregation Group can connect Portals in networks that are under separate administration, running different fault recovery protocols.
- e) **Administrative independence**—The specification of DRNI to interconnect separate networks does not introduce new requirements on either networks' existing control protocols.
- f) **Inter-network fault isolation**—The failure or recovery of a link or node in one network, requiring a reaction by that network's control protocols, can be hidden by DRNI from the second network's control protocols. Thus, super-networks can be created out of separate networks interconnected via DRNI, without propagating one network's fault and recovery events throughout the super-network.
- g) **Network-DRNI fault isolation**—The failure or recovery of a link between two Portals can be hidden by DRNI from both networks' control protocols.
- h) **Rapid fault recovery**—Means for the Systems in a Portal to communicate are provided so that they can cooperate to respond rapidly to failure or recovery events, typically on the order of milliseconds for link down events and 1 second or less for link up events.
- i) **Extended faults**—Optional elements of DRNI can support three Systems in a Portal, so that fault redundancy can be provided even while a System is added or removed from a Portal.
- j) **Distribution independence**—The frame distribution algorithm used to satisfy network requirements can be different from the algorithm used to assign frames to the Aggregation Ports of a Link Aggregation Group.


The Distributed Resilient Network Interconnect, as specified in this Clause, does not support the following:

- 1 k) **Multipoint Aggregations**—Connections among more than two Portals or Aggregation Systems are
2 not supported.
- 3 l) **Virtual System emulation**—DRNI makes no attempt to provide all of the mechanisms required to
4 make two Systems, such as two Bridges or two Routers, appear, to all other Systems, even those in
5 their own network, to be a single (virtual) System. But, neither does DRNI prevent the Systems
6 employing it to emulate a single virtual System.
- 7 m) **More than three Systems in a Portal**—In general, a fault-tolerant Portal comprising more than
8 three Systems would require some protocol mechanism, not provided by this standard, to select
9 forwarding paths and to prevent endless forwarding loops.

9.2 Distributed Relay

14 DRNI is created by using a Distributed Relay to interconnect two or three Systems, each running Link
15 Aggregation, to create a Portal. Each System in the Portal (i.e., each Portal System) runs Link Aggregation
16 with a single Aggregator. The Distributed Relay enables the Portal Systems to jointly terminate a Link
17 Aggregation Group. To all other Systems to which the Portal is connected, the Link Aggregation Group
18 appears to terminate in a separate emulated System created by the Portal Systems.

19 Figure 9-1 illustrates the starting point for describing the Distributed Relay. The (three) network links
20 depicted in the Figure 9-1 and discussed in this standard correspond to physical or logical links that are
21 visible to and are under the control of network protocols. In this diagram, Systems **A** and **B** each are
22 characterized by performing a “Function 1,” which is presumably some kind of packet relay function, e.g., a
23 router or a bridge. “Function 1” could just as well be a file server operation, in which case the outside two
24 “ports” on each System would likely not be present. Each System runs a single instance of a Link
25 Aggregation sublayer. Let us suppose that it is desired that the shaded ports be associated into a Portal with a
26 Distributed Relay.



38 **Figure 9-1—Distributed Relay: starting point**

39 Figure 9-1 is an example, not the general case. In general, the Distributed Relay supports:

- 40
- 41 a) The necessary protocols and procedures for only the configurations listed in 9.3.1 as provided by
42 this standard.
- 43 b) Link Aggregation functions, each subsuming one or more MACs
- 44 c) Connections among the Portal Systems of the Distributed Relay.

45 The purpose of introducing the Distributed Relay functional layer in this example is to make these two
46 Portal Systems appear, to the Systems connected to them, to be in the configuration illustrated in Figure 9-2.
47 There appears to exist a third emulated System **C**, connected to the original Portal Systems by a link that has
48 been inserted between Function 1 and Link Aggregation. That is, Portal Systems **A** and **B** conspire to
49 behave, insofar as any other Systems to which they are connected can discern, as if emulated System **C**
50 actually exists, as shown in Figure 9-2. While Figure 9-2 is an example, it illustrates the principles of the
51 Distributed Relay:

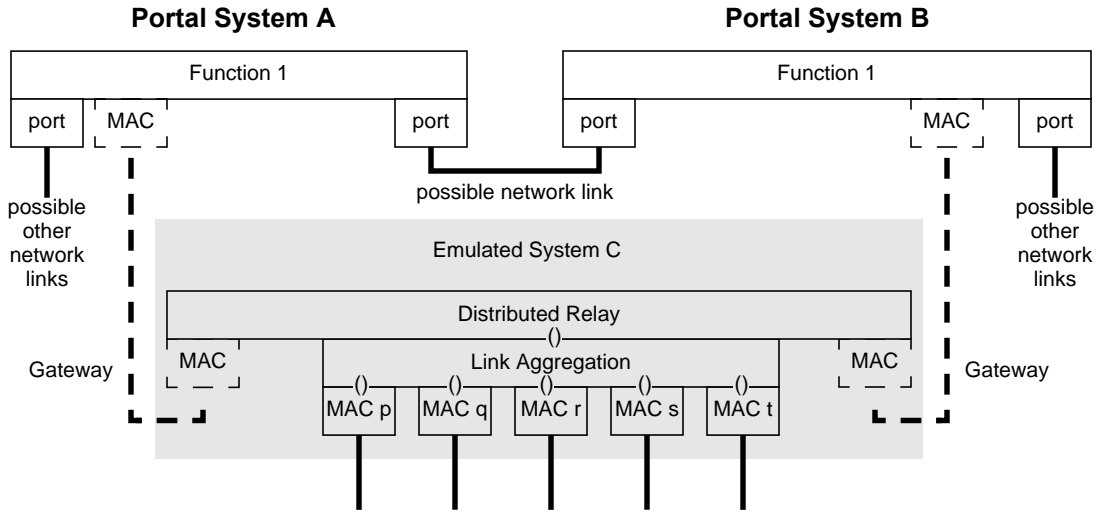


Figure 9-2—Distributed Relay: as seen by other Systems

- d) The Distributed Relay in the emulated System C is an (N+1)-port relay for N Portal Systems, with N Gateway Ports connected to the Portal Systems, and a single Emulated Link Aggregation sublayer associated with the original Portal Systems.
- e) The Aggregation Ports (MACs) have been moved to the emulated System, and thus appear, to all other Systems, to be equally distant from the real Portal Systems comprising the Distributed Relay.

The actual constructs used by the two Portal Systems to create the emulated System C are illustrated in Figure 9-3. Figure 9-3 shows the two DR Functions of the Distributed Relay, one DR Function in each System A and B. This example illustrates the remaining principles of the Distributed Relay:

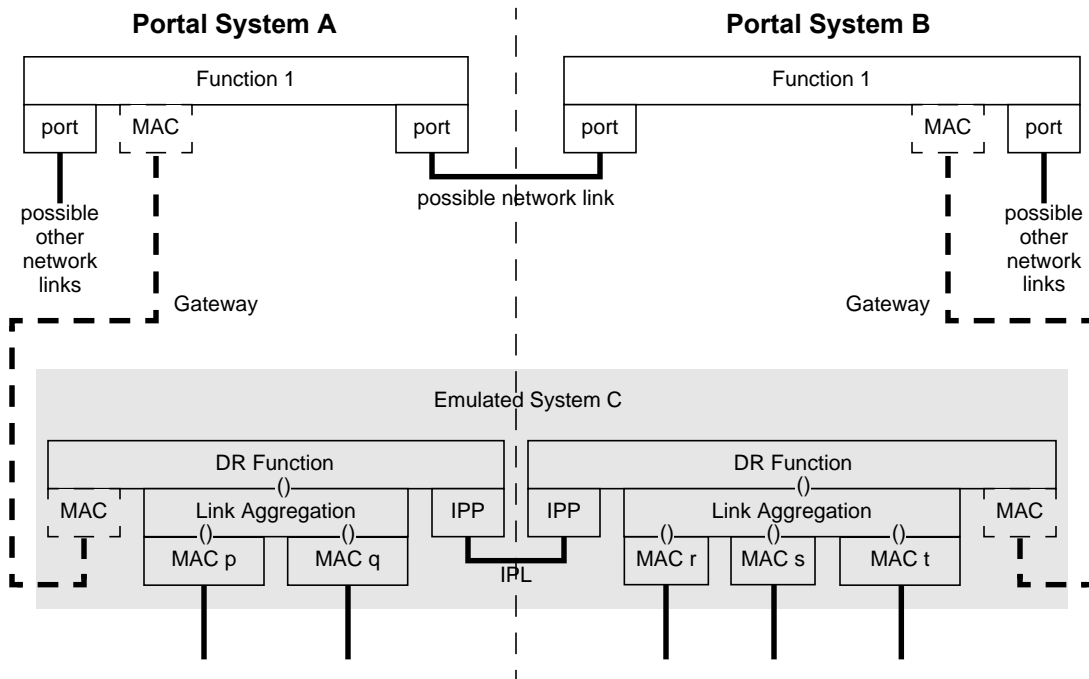
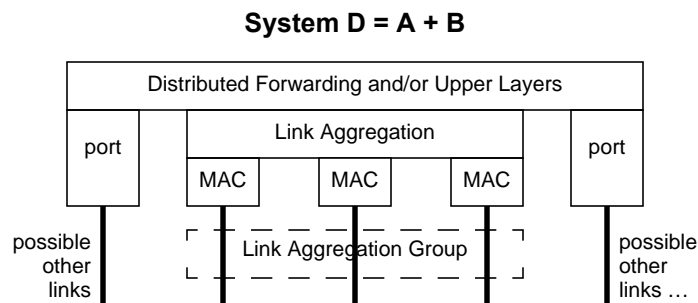


Figure 9-3—Distributed Relay: as seen by Systems A and B

- 1 f) In each System **A** and **B**, the ports that are to be associated with System **C** are moved to a position
2 below the DR Function's Link Aggregation sublayer.
3 g) A virtual link and its terminating virtual MACs, called a "Gateway", is constructed to connect each
4 DR Function to its Function 1.
5 h) Between each pair of DR Functions in the Portal there is constructed an Intra-Portal Link (IPL),
6 terminated at each end by an Intra-Portal Port (IPP). (This can exist in many forms; see 9.3.2.)
7 i) There is a "Gateway algorithm" that decides through which Gateway a frame can pass into or out of
8 the emulated Distributed Relay.
9 j) Similarly, a "Port algorithm" decides through which Portal System's Aggregation Port a frame can
10 pass into or out of the emulated Distributed Relay.
11 k) As mentioned above, there can be three Systems participating to create a Portal and a emulated
12 System **C**. In that case, the Distributed Relay in emulated System **C** has an additional Gateway Port,
13 one to each Portal System, and two or three IPLs to interconnect the DR Functions.
14 l) The DR Functions, as specified in 9.3, work together to move frames between the Gateways, the
15 IPL, and the Link Aggregation sublayers.
16

17 NOTE 1—This Standard does *not* define the alternate model shown in Figure 9-4, in which the entirety of Systems **A**
18 and **B** simulate a single System **D**, but neither does this Standard prevent the use of DRNI in such a model.
19



21
22
23
24
25
26
27
28
29 **Figure 9-4—Not an example of a DRNI Portal**

30
31
32 NOTE 2—The support of Figure 9-2 is important for achieving goal item e) in 9.1, administrative independence, without
33 excessive standardization. System **C** in Figure 9-2 can be made visible to the network control protocols. This makes it
34 relatively easy for Systems **A** and **B** to hide events occurring in the Link Aggregation Group or in the other Portal from
35 the other Systems in their own network, because System **C** is connected only to Systems **A** and **B** in that network.
36 Systems **A** and **B** maintain their own separate identities in their own network. They conspire to behave, insofar as any
37 other Portals or Systems to which they are connected can discern, as if emulated System **C** as shown in Figure 9-2
38 actually exists, but if the Portals are in separately administered networks, this task is relatively easy. In the scheme
39 shown in Figure 9-4 on the other hand, Systems **A** and **B** have to appear to be a single System **D** to all other Systems in
40 their own network. This is, in general, a much more difficult task to accomplish, and one not easily standardized.
41
42

43 9.3 Distributed Relay operation and procedures

44
45 The DR Function in each Portal System (Figure 9-3) is intended to have (subject to operational failures)
46 three kinds of ports:
47

- 48 a) Intra-Portal Ports (9.3.2), at most one (perhaps complex; see 9.3.2) IPL Port connected to each of the
49 other Portal System(s) belonging to the same Portal;
50 b) Exactly one virtual Gateway Port with a virtual Link to a virtual Gateway Port in the Portal System
51 in which the DR Function resides; and
52 c) Exactly one Aggregator Port (the port that is supported by the ISS instance identified by the prefix
53 *Agg*: in 6.2.2) to the Link Aggregation sublayer with any number of Aggregation Ports, intended to
54

1 connect to other Systems in a manner such that those other Systems believe they are connected to a
2 single emulated System.

3
4 In Figure 9-2, the IPLs and IPPs are not visible, and the Distributed Relay of the emulated Aggregation
5 System **C** has one Gateway to each of the Systems in its Portal.

6
7 The purpose of the emulated Distributed Relay is to pass every frame received from an Aggregator Port (“up
8 frame”) to a single Gateway, or to discard it, and every frame received from a Gateway (“down frame”) to an
9 Aggregator Port, or discard it. The DR Functions comprising the Distributed Relay sometimes has to send a
10 frame across one or two IPLs in order to get it to the correct Gateway or Aggregator Port. A DR Function
11 makes the choice of whether to discard or pass a frame to its Gateway, Aggregator Port, or one of its IPLs by
12 assigning every frame to two Conversation IDs, a Gateway Conversation ID and a Port Conversation ID,
13 and configuring the Gateways, Aggregation Ports, and IPLs in terms of these Conversation IDs.

14
15 The “Gateway algorithm” mentioned in item i) in 9.2 consists of two parts, an algorithm for assigning any
16 given frame to a Gateway Conversation ID (7.4.1.1.13), and an assignment of Gateway Conversation IDs to
17 Gateways (Drni_Gateway_Conversation, 9.3.4.2).

18
19 NOTE 1—If the Portal System is a VLAN Bridge performing learning, the mapping of a frame to a Gateway
20 Conversation ID will be based on its VLAN ID to avoid oscillations in the learning process within the attached network.
21 For implementations of DRNI in these cases there can be a one-to-one map of VLAN ID to Gateway Conversation ID.

22
23 Similarly, the “Port algorithm” of item j) in 9.2 consists of an algorithm for assigning any given frame to a
24 Port Conversation ID (e.g., 8.2), and an assignment of Port Conversation IDs to Aggregation Ports
25 (Drni_Port_Conversation, 9.3.4.2). The Port algorithm is identified by the aAggPortAlgorithm (7.3.1.1.33)
26 and is also used in the operation of Conversation-sensitive frame collection and distribution described in 6.6.

27
28 Means are specified in 9.4 to ensure that all of the DR Functions on a given Portal use the same Gateway
29 algorithm and the same Port algorithm for assigning frames to their respective Conversation IDs, and to
30 guarantee that any given Gateway Conversation ID is assigned to at most one of the Gateways, and any
31 given Port Conversation ID to at most one of the Aggregation Ports of a Portal, at any given moment.

32
33 It is allowed, but not required, that the Gateway algorithm and the Port algorithm use the same means for
34 assigning a frame to a Conversation ID, so that the Gateway Conversation ID equals the Port
35 Conversation ID.

36
37 NOTE 2—The Gateway selection process is in general dependent on the service parameters and forwarding capabilities
38 of the hosting System while the Port selection process is dependent on the interface capabilities. For example a Bridge
39 might use a Gateway algorithm based on VIDs while using a Port algorithm based on I-SIDs for its I-tagged interfaces.

40
41 The Port algorithm is always applied to a frame as it enters the DR Function from the Gateway to determine
42 whether to send it to the Aggregator Port or to a particular IPP. The Gateway algorithm is always applied to
43 a frame as it enters from the Aggregator Port to determine whether to send it to the Gateway or to a
44
45
46
47
48
49
50
51
52
53
54

particular IPP. Both algorithms have to be applied, and their results compared, in order to forward a frame entering a DR Function from an IPL, as shown in Table 9-1.

Table 9-1—DR Function: forwarding frame received from IPL n

Gateway algorithm says, “Gateway is ...”	Port algorithm says, “Aggregation Port is ...”	DR Function emits frame to:
my Gateway	any ^a	my Gateway
behind IPL n	one of my Aggregation Ports	my Aggregator Port
	behind IPL n	discard ^b
	behind IPL $m (\neq n)$	IPL $m (\neq n)$
behind IPL $m (\neq n)$	any ^a	IPL $m (\neq n)$

- a “Any” means that the output from the Port algorithm is not used; the Gateway algorithm determines to which port the frame is sent.
- b DR Functions in different Portal Systems have incompatible configurations, or there is a malfunction. Discard frame to prevent looping.

The rationale for this table is as follows:

- d) Table 9-1 assumes one of three configurations (see 9.3.1):
- 1) Two Portal Systems connected by a single IPL;
 - 2) Three Portal Systems connected linearly by two IPLs; or
 - 3) Three Portal Systems connected circularly by three IPLs.
- e) The Gateway algorithm is enforced in both directions through the Gateway; that is, a frame entering the DR Function from a Gateway is discarded if the Gateway algorithm, applied to that frame, would not send it back up the Gateway. This is necessary to prevent frames received from a Gateway being forwarded across an IPL and passed back into the network through another Gateway.
- f) If the Gateway algorithm indicates that the frame should pass through the Gateway, it has to be an up frame, because a down frame could not enter the Portal from any other DR Function, by item e.
- g) Otherwise, if the Gateway algorithm indicates that the frame came from the IPL to which it would be forwarded, then it is a down frame, and so is forwarded using the Port algorithm. (If the Port algorithm would send it back on the port on which it arrived, there is some sort of malfunction or misconfiguration, and the frame is discarded.)
- h) Otherwise, the Gateway algorithm indicates that the frame did not come from the IPL to which it would be forwarded, so it has to be an up frame, and the frame is directed according to the Gateway algorithm.

NOTE—The Port algorithm and Distributed Relay Control Protocol of the Distributed Relay together determine the mapping of Port Conversation IDs to individual Aggregation Ports, and the variables in their control determine the operation of the Frame Distributor and Frame Collector. However, this does not alter the path of data and control as described in Clause 6, since the Distributed Relay passes all data to or from the Aggregation Ports through the Aggregator.

The application of the Gateway and Port algorithms on frames entering a DR Function from a Gateway and an Aggregator Port, as shown in Table 9-2 and Table 9-3 respectively.

Table 9-2—DR Function: forwarding frame received from my Gateway

Gateway algorithm says, “Gateway is ...”	Port algorithm says, “Aggregation Port is ...”	DR Function emits frame to:
my Gateway	one of my Aggregation Ports	my Aggregator Port
	behind IPL <i>n</i>	IPL <i>n</i>
behind IPL <i>n</i>	any	discard

Table 9-3—DR Function: forwarding frame received from one of my Aggregation Ports

Gateway algorithm says, “Gateway is ...”	Port algorithm says, “Aggregation Port is ...”	DR Function emits frame to:
my Gateway	any	my Gateway
behind IPL <i>n</i>	any	IPL <i>n</i>

9.3.1 Portal Topology

The most general topology of a Portal could be

- a) three Portal Systems connected in a ring by three Intra-Portal Links.

Other supported topologies are subsets of this, including:

- b) three Portal Systems connected in a chain by two IPLs,
- c) two Portal Systems connected by a single IPL,
- d) a Portal System with no active IPLs.




Figure 9-5—Portal Topology

These subsets can arise by design or as a result of system or link failures.

The terms Home, Neighbor, and Other neighbor are used to identify the Portal Systems from the point of view of a given Intra-Portal Port. The Home is the system containing the IPP. The Neighbor is the system connected to the IPP. The Other neighbor is the system connected to the other IPP (if present) in the Home system. Using IPP B1 for an example, its home system is B, its immediate neighbor is A (because it is connected to IPP B1 via IPL AB), and its other neighbor is C (because it is connected to IPP B2 via IPL BC).

1 BC). Note that the Other neighbor of IPP A2 is also C (because it is connected to IPP A1 via IPL AC), so
2 comparing the IDs of the Other neighbors of IPPs B1 and A2 verifies that there are no more than three
3 systems in a ring or chain.
4

5 **9.3.2 Intra-Portal Link**

6
7 An Intra-Portal Link (IPL) is a single logical point-to-point link between DR Functions in two different
8 Systems. A DR Function has at most one IPL for each of the other Systems comprising one end of a DRNI
9 Link Aggregation Group. IPLs and network links can share a physical link or link aggregation.

10
11 An IPL can be physical (e.g., an 802.3 Ethernet LAN) or logical (e.g., a 802.1Q Backbone Service Instance
12 or an IETF pseudowire). An Intra-Portal Link can share a physical link with other Intra-Portal Links or
13 network links. An Intra-Portal Link can be a Link Aggregation Group, and thus consist of a number of
14 physical links.
15

16 It will often be the case in deployed networks that two Systems will be configured with both a normal
17 network link connecting them, and an IPL, as illustrated in Figure 9-3. It would reduce the utility of DRNI if
18 every Portal required its own separate physical IPL, especially if a pair of Systems are configured to support
19 multiple Portals. DRNI supports a number of methods by which the Systems can distinguish frames on a
20 network link from frames on a particular IPL:
21

- 22 a) **Physical.** A separate physical link can be used to support any particular network link or IPL.
- 23 b) **Aggregated.** A separate Aggregator Port can be used for supporting an IPL.
- 24 c) **Time-shared.** A network link and one or more IPLs can use the same physical link (or Aggregator
25 Port), but at different times. This requires that the Systems disable the use of the network link when
26 the IPL is required for connectivity, or else that the use of the Aggregation Links and the selection of
27 Gateways be adjusted to eliminate the need for using the IPL when the network link is required. This
28 technique is described in 9.3.2.1.
- 29 d) **Tag-shared.** A network link and one or more IPLs can use the same physical link (or Aggregator
30 Port), using different Service IDs. Tag sharing is described in 9.3.2.2.
- 31 e) **Logical.** The frames on the network link and the IPL(s) can be encapsulated, as described in 9.3.2.3.
32

33 A System implementing the DRNI shall support using separate physical links for IPLs and network links,
34 and may support any of the other methods.
35

36 At each end of the shared physical link or Aggregator Port, there is one virtual port for each function
37 (network link or a particular IPL) that the link or Aggregator Port is used for. Any of the above methods can
38 be used simultaneously between two Systems, by administrative choice, as long as both ends of any given
39 physical link or Aggregator Port use the same method.
40

41 **9.3.2.1 Network / IPL sharing by time**

42
43 The goal of Network/IPL sharing by time is to support DRNI without requiring separate physical links for a
44 network connection and the IPL, and without requiring any frame modification. When sharing a link, it is
45 necessary to be able to determine for each frame whether that frame is intended to be traversing the network
46 connection or the IPL. This determination can be made without modifying the frame (e.g. without
47 translating the VLAN ID or adding a tag or encapsulation) if at any given time the physical link is known to
48 be only used as a network link or only used as an IPL. Whether the link is used as a network link or an IPL
49 at any given time is established by the control protocol used to establish a fully-connected, loop-free active
50 topology for each VLAN in the network.
51

- 52 a) If the link is not included in the active topology of a VLAN (e.g. it has been blocked by the
53 operation of the network control protocol), it is available to be used as the IPL. In this case the link
54 is used by DRNI just as if it were a dedicated (unshared) IPL.

- 1 b) If the link is included in the active topology of a VLAN, then there is no IPL available for that
2 VLAN. In this case the DRNI is unable to pass a frame between an Aggregation Port in one Portal
3 System and a Gateway in another Portal System. Therefore for any given frame, DRNI is restricted
4 to have the Gateway and the Aggregation Port in the same Portal System.
5

6 NOTE 1—The fact that the shared link can be unavailable for the transfer of frames between a Gateway and a specific
7 Aggregation Port does not restrict the ability to exchange DRCPDUs on the shared link.
8

9 There are two cases to consider in meeting the restriction that for any given frame the Gateway and the
10 Aggregation Port are in the same Portal System. The straightforward case is when the Port Conversation IDs
11 are agreed and symmetric across the DRNI, and all frames with a given VLAN ID map to the Port
12 Conversation IDs that select Aggregation Ports in the same Portal System. Then that Portal System is
13 selected as the Gateway for that VLAN ID, and there is no need for any data frames to traverse the IPL. In
14 any other circumstance the only way to assure that the Gateway and a specific Aggregation Port are in the
15 same Portal System is that when a Portal System receives a frame on any port other than the shared network/
16 IPL port, that Portal System is considered the Gateway for that frame and in each Portal System all Port
17 Conversation IDs map to an Aggregation Port attached to that System. In this mode, a Portal System which
18 receives any frame on a network port (not being the IPP or the Aggregation Port) is responsible for
19 forwarding that frame to the Aggregation Port if required. A Portal System which receives a frame on the
20 IPP never forwards that frame to the Aggregation Port. In this case the gateway selection cannot necessarily
21 be based on VID, so when the Portal Systems are 802.1Q Bridges the learning process on the shared
22 network/IPL link is compromised. Because the learning issues are limited to this port, it can be remedied by
23 synchronizing the addresses learned on the DRNI with the other Portal System. See Annex G for further
24 details on MAC Address Synchronization.
25

26 **9.3.2.2 Network / IPL sharing by tag**

27
28 If per-service frame distribution (8.2) is employed, and if the number of services required to support the
29 network link, plus the number of services required to support one or more IPLs, is less than the number of
30 services supplied by the frame format used (e.g., 4094 S-VLAN IDs), then VID translation can be used to
31 separate the frames on the different logical links.
32

33 The method is selected by configuring the `aDrniEncapsulationMethod` (7.4.1.1.17) with the value 2. If
34 enabled, as indicated by the variable `Enabled_EncTag_Shared` which is controlled by the Network/IPL
35 sharing machine (9.4.20), every frame which is to be transported by the IPL to the Neighbor Portal System
36 and is associated with a Gateway Conversation ID, is translated to use a value configured in
37 `aDrniIPLEncapMap` (7.4.1.1.18) and every frame, that is to be transported by network link, shared with the
38 IPL, and is associated with a Gateway Conversation ID, is translated to use a value configured in
39 `aDrniNetEncapMap` (7.4.1.1.19).
40

41 **9.3.2.3 Network / IPL sharing by encapsulation**

42
43 This method enables sharing an IPL with a network link by using encapsulation techniques (e.g., a 802.1Q
44 Backbone Service Instance, a B-VLAN, an IETF pseudowire, etc.)
45

46 The method is selected by configuring the `aDrniEncapsulationMethod` (7.4.1.1.17) with a value representing
47 the encapsulation method that is used to transport IPL frames to the Neighbor Portal System when the IPL
48 and network link are sharing the same physical link. This value consists of the three-octet OUI or Company
49 Identifier (CID) identifying the organization which is responsible for this encapsulation and one following
50 octet used to identify the encapsulation method defined by that organization. If enabled, as indicated by the
51 variable `Enabled_EncTag_Shared` which is controlled by the Network/IPL sharing machine (9.4.20), every
52 frame, that is to be transmitted on the IPL to the Neighbor Portal System and is associated with a Gateway
53 Conversation ID, is encapsulated with a method specified by `aDrniEncapsulationMethod` to use a value
54

1 configured in aDrniIPLEncapMap (7.4.1.1.18) and every frame, that is received by the IPL is de-
2 encapsulated and mapped to a Gateway Conversation ID using the same table.
3

4 **9.3.3 Protocol Identification**

5
6 All Distributed Resilient Network Interconnect protocols use the Distributed Resilient Network Interconnect
7 EtherType value as the primary means of protocol identification; its value is shown in Table 9-4.
8

9
10 **Table 9-4—DRNI EtherType Assignment**

11 Assignment	12 Value
13 Distributed Resilient Network Interconnect (DRNI) 14 EtherType	15 89-52

16
17
18 The first octet of the MAC Client data following the EtherType field is a protocol subtype identifier that
19 distinguishes between different DRNI protocols. Table 9-5 identifies the semantics of this subtype.
20

21
22 **Table 9-5—DRNI Protocol subtypes**

23 Protocol Subtype value	24 Protocol name
25 0	26 Reserved for future use
27 1	28 Distributed Relay Control Protocol (DRCP) (9.4)
29 2	30 Address Synchronization Protocol(Annex G)
31 3-255	32 Reserved for future use

33 **9.3.4 DR Function state machines**

34
35 The DR Function state machines shall implement the forwarding rules specified in Table 9-1, Table 9-2, and
36 Table 9-3. These forwarding rules can be summarized as follows:
37

- 38 a) For frames entering through an Aggregation Port, Up frames, the Gateway algorithm decides
39 whether to transmit it to the Gateway link or to the IPP according to its Gateway Conversation ID. If
40 the frame's Gateway Conversation ID matches the Portal System's operational Gateway
41 Conversation ID, the frame will be forwarded to the Gateway, otherwise it will be forwarded to the
42 IPP.
- 43 b) For frames entering through a Gateway, Down frames, the Port algorithm decides whether to
44 transmit it to the Aggregation Port or to the IPP according to its Port Conversation ID. If the frame's
45 Port Conversation ID matches the Portal System's operational Port Conversation ID, the frame will
46 be forwarded to the Aggregation Port, otherwise it will be forwarded to the IPP.
- 47 c) An Up frame offered to an IPP is only transmitted if the Portal System for this Gateway
48 Conversation ID lies behind that IPP, otherwise it is discarded.
- 49 d) A Down frame offered to an IPP is only transmitted if the target Portal System for this Port
50 Conversation ID lies behind that IPP, otherwise it is discarded
51

52
53 Some of the Link Aggregation variables used by a Distributed Relay have to be formed in a particular
54 manner, so that multiple Portal Systems can cooperate to create a single emulated System:

- 1 e) The two least significant bits of the Port Priority in the Port ID (6.3.4) for each Aggregation Port in a
2 Distributed Relay's Aggregator Port are set to the value of DRF_Portal_System_Number.
3 f) The most significant two bits of the Administrative Key for each Aggregation Port and associated
4 Aggregator in a Distributed Relay's Aggregator Port are set to the value of
5 DRF_Portal_System_Number. The remainder of the bits can be used as described in 6.3.5 to reflect
6 the physical characteristics of the Aggregation Ports and they should be set consistently across the
7 Portal Systems in a Portal so that a single Aggregator can be formed across the Distributed Relay.
8

9 The DR Function runs Link Aggregation with a single Aggregator. If the DR Function's Aggregation Ports
10 can form more than one Link Aggregation Group, only one LAG can be attached to the DR Function's
11 Aggregator (6.4.14.1).
12

13 9.3.4.1 Service interfaces

14 Since a DR Function uses various instances of the ISS, it is necessary to introduce a notation convention so
15 that the reader can be clear as to which interface is being referred to at any given time. A prefix is therefore
16 assigned to each service primitive, indicating which of the interfaces is being invoked. The prefixes are as
17 follows:
18

- 19 a) *Agg.*, for primitives issued on the interface between the DR Function and the Link Aggregation
20 sublayer.
21 b) *Gate.*, for primitives issued on the Gateway.
22 c) *MacIppN.*, for primitives issued on the interface between the MAC entities supporting the IPL *n* and
23 the DRCP Control Parser/Multiplexer (9.4.4).
24 d) *DRCPCtrlMuxN.*, for primitives issued on the interface between DRCP Control Parser/Multiplexer
25 N and the DRCP control entity (where N is identifying the IPP associated with the DRCP Control
26 Parser/Multiplexer).
27 e) *IppN.*, for primitives issued on the interface of the DR Function supported by the DRCP Control
28 Parser/Multiplexer N (where N is identifying the IPP associated with the DRCP Control Parser/
29 Multiplexer).
30
31

32 9.3.4.2 Per-DR Function variables

33 DA

34 SA

35 mac_service_data_unit

36 priority

37 The parameters of the M_UNITDATA.indication primitive.

38 BEGIN

39 A Boolean variable that is set to TRUE when the System is initialized or re-initialized, and is
40 set to FALSE when (re-)initialization has completed.

41 Value: Boolean

42 Drni_Portal_System_Gateway_Conversation

43 Operational Boolean vector, indexed by Gateway Conversation ID, indicating whether the
44 indexed Gateway Conversation ID is allowed to pass through this DR Function's Gateway
45 (TRUE = passes). Its values are computed by the updatePortalSystemGatewayConversation
46 function (9.4.11).
47

48 Value: sequence of Boolean values, indexed by Gateway Conversation ID.

49 Drni_Portal_System_Port_Conversation

50 Operational Boolean vector, indexed by Port Conversation ID, indicating whether the indexed
51 Port Conversation ID is allowed to be distributed through this DR Function's Aggregator
52 (TRUE = passes). Its values are computed by the updatePortalSystemPortConversation
53 function (9.4.11).
54

Value: sequence of Boolean values, indexed by Port Conversation ID.

9.3.4.3 Per-IPP variables

Ipp_Gateway_Conversation_Direction

Operational Boolean vector of which Gateway Conversation IDs are assigned to Gateways reachable through this IPP. It is set by the operation of DRCP.

Value: Vector of Boolean flags indexed by Gateway Conversation ID; TRUE = some Gateway reachable through this IPP is enabled for this Gateway Conversation ID.

Ipp_Gateway_Conversation_Direction is initialized to FALSE and recomputed by the updateIPPGatewayConversationDirection function (9.4.11) whenever any of its contributing variables changes. For frames received on this IPP, TRUE means that the frame is a Down frame, ultimately destined for an Aggregator (or discard), and FALSE means the frame is an Up frame, ultimately destined for a Gateway (or discard). For frames offered for transmission on this IPP, TRUE indicates that the frame can pass, and FALSE that it cannot.

Ipp_Port_Conversation_Passes

Operational Boolean vector of which Port Conversation IDs are allowed to be transmitted through this IPP.

Value: Vector of Boolean flags indexed by Port Conversation ID.

This variable is examined only when a Down frame is offered for transmission on this IPP.

Ipp_Port_Conversation_Passes is initialized to FALSE and recomputed by the updateIPPPortConversationPasses function (9.4.11) whenever any of its contributing variables changes.

9.3.4.4 Functions

extractGatewayConversationID

This function extracts a Gateway Conversation ID value by applying the Gateway Algorithm (7.4.1.1.13) to the values of the parameters of the service primitive that is invoked on the DR Function's Relay entity on receipt of an ISS primitive at one of the DR Function's port. The relationship of the parameter values on the ISS primitives and the service primitives on the DR Function's Relay entity ports is provided by the associated supporting functions on those ports and their configuration.

NOTE—These supporting functions can be as simple as the EISS supporting functions specified in 6.9 of IEEE Std 802.1Q-2011, for the case of a DRNI supported on a Customer Network Port or a Provider Network Port on a Provider Bridge (Clause 15 in IEEE Std 802.1Q), or more complex, like the EISS supporting functions specified in 6.10 or 6.11 in IEEE Std 802.1Q-2011, for the case of DRNI supported on a Provider Instance Port or a Customer Backbone Port respectively on a Backbone Edge Bridge (Clause 16 in IEEE Std 802.1Q) or, like the C-tagged Service Interface supporting functions or the Remote Customer Service Interface supporting functions specified in 15.4 or 15.6 in IEEE Std 802.1Q-2013 for the case of DRNI supported on a Customer Edge Port or a Remote Customer Access Port respectively on a Provider Edge Bridge.

Value: Integer in the range of 0 through 4095.

extractPortConversationID

This function extracts a Port Conversation ID value by applying the Port Algorithm (7.3.1.1.33) to the values of the parameters of the service primitive that is invoked on the Aggregator on receipt of a ISS primitive at one of the other DR Function's port. The relationship of the parameter values on the ISS primitives on the Aggregator and the corresponding service primitives on the DR Function's port is provided by the associated supporting function on the Aggregator and the DR Function port and their configurations. Check the NOTE above.

Value: Integer in the range of 0 through 4095.

9.3.4.5 Messages

- Agg:M_UNITDATA.indication
- Gate:M_UNITDATA.indication
- IppN:M_UNITDATA.indication
- Agg:M_UNITDATA.request
- Gate:M_UNITDATA.request
- IppN:M_UNITDATA.request

The service primitives used to pass a received frame to a client with the specified parameters.

If Network / IPL sharing by tag (9.3.2.2), or Network / IPL sharing by encapsulation (9.3.2.3) methods are used, the service primitives IppN:M_UNITDATA.indication and IppN:M_UNITDATA.request need to be manipulated through the operation of functions which are controlled by the Network/IPL sharing machine in 9.4.20.

9.3.4.6 DR Function: Aggregator Port reception state machine

The DR Function: Aggregator Port reception state machine shall implement the function specified in Figure 9-6 with its associated parameters (9.3.4.2 through 9.3.4.5). There is one DR Function: Aggregator Port reception state machine per Portal System and there are as many PASS_TO_IPP_N states as IPPs in a Portal System, each identified by the index *n*. The prefix “n.” in the diagram is used to identify the specific IPP *n* that the associated variable is related to.

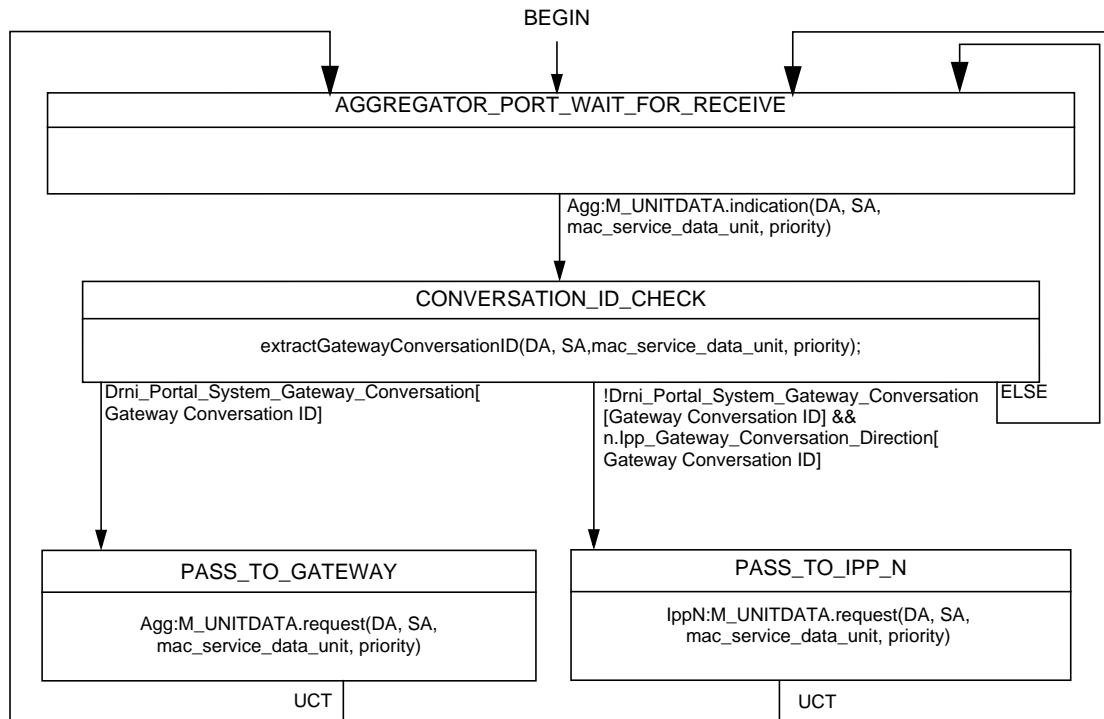


Figure 9-6—DR Function: Aggregator Port reception state machine

9.3.4.7 DR Function: Gateway distribution state machine

The DR Function: Gateway distribution state machine shall implement the function specified in Figure 9-7 with its associated parameters (9.3.4.2 through 9.3.4.5). There is one DR Function: Gateway distribution

state machine per Portal System and there are as many PASS_TO_IPP_N states as IPPs in a Portal System, each identified by the index *n*. The prefix “n.” in the diagram is used to identify the specific IPP *n* that the associated variable is related to.

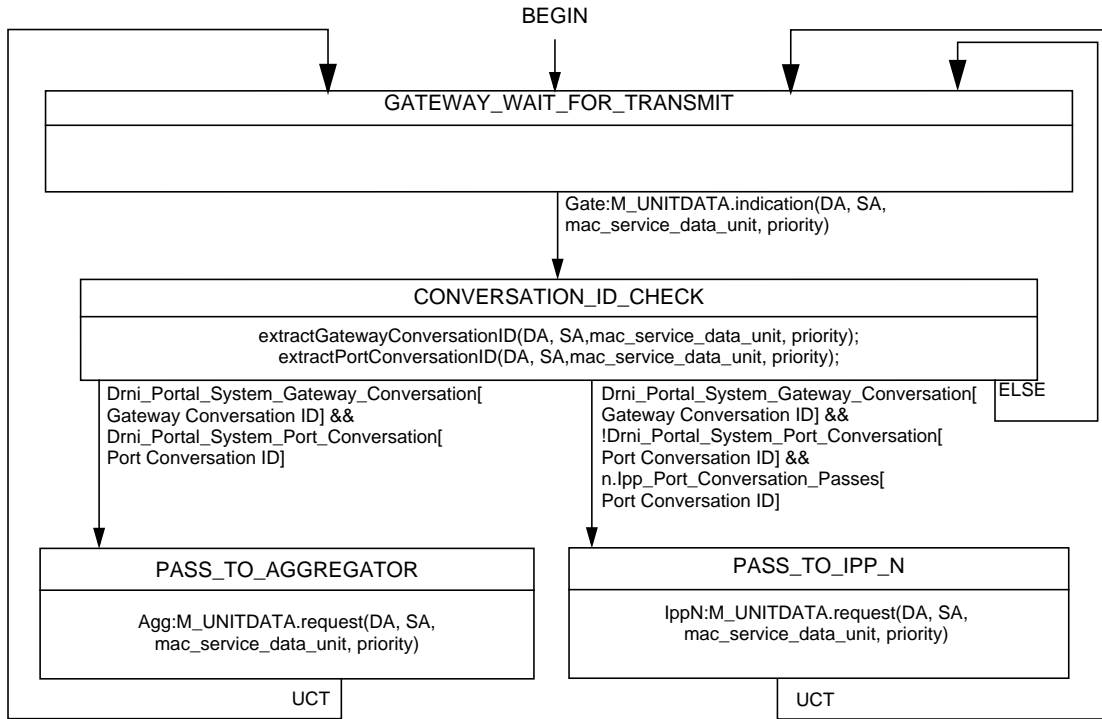


Figure 9-7—DR Function: Gateway distribution state machine

9.3.4.8 DR Function: IPP N reception state machine

The DR Function: IPP N reception state machine shall implement the function specified in Figure 9-8 with its associated parameters (9.3.4.2 through 9.3.4.5). There is one DR Function: IPP N reception state machine per IPP per Portal System and there are as many PASS_TO_IPP_M states as IPPs in a Portal System, each

identified by the index m . The prefix “n.” or “m.” in the diagram is used to identify the specific IPP n or IPP m that the associated variable is related to.

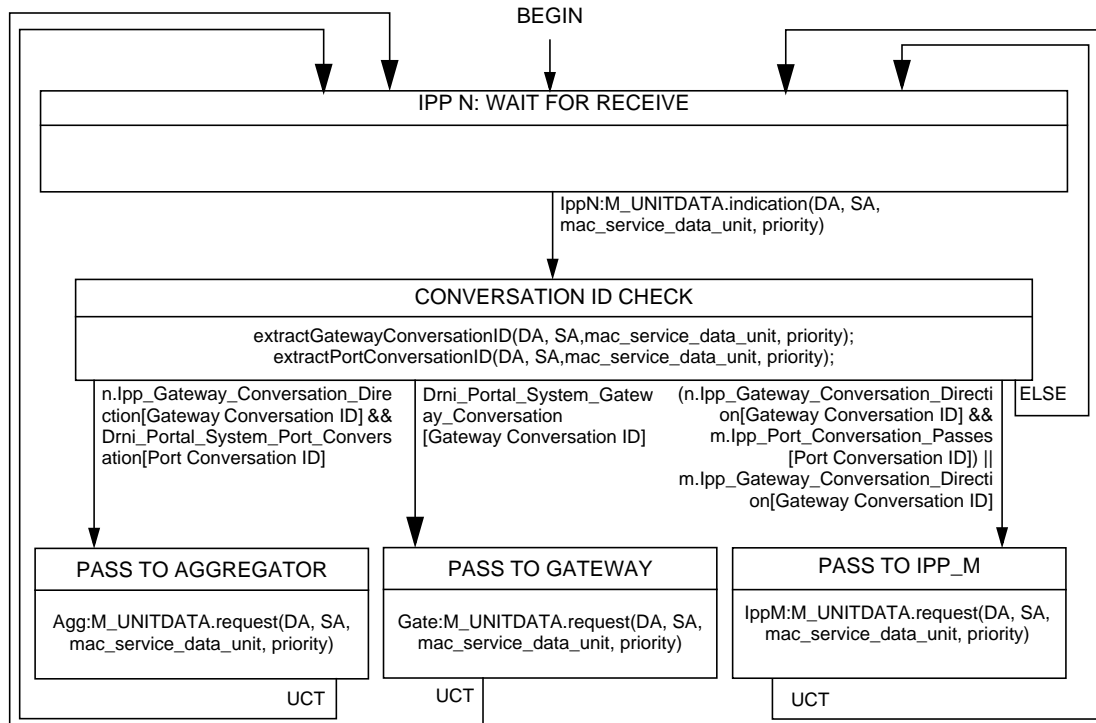


Figure 9-8—DR Function: IPP N reception state machine

9.4 Distributed Relay Control Protocol

The purpose of the Distributed Relay Control Protocol (DRCP) is to:

- a) Establish communication between Portal Systems across an Intra-Portal Link (9.4.1);
- b) Verify the consistent configuration of Portal Systems (9.4.1);
- c) Determine the identity to be used for the emulated System;
- d) Distribute the current states of the Portal Systems and their Aggregation Ports among each other;
- e) Compute the resultant path of any frame required to pass through each IPL, and exchange the information with adjacent Portal Systems as required to ensure against forwarding loops and duplicate frame delivery.
- f) Maintain Link Aggregation with a single Aggregator across the Portal.

The result of the operation of DRCP is to maintain the variables (9.3.4.2, 9.3.4.3) that control the forwarding of frames by the Distributed Relay.

DRCP exchanges information to ensure that the Portal Systems can work together. The first class of such information includes the managed objects and variables that have to be compatible in order to pass any data at all (item a). These include:

- g) **aAggPortAlgorithm:** All Portal Systems have to use the same Port algorithm.
- h) **aDrniGatewayAlgorithm:** All Portal Systems have to use the same Gateway algorithm.
- i) **aDrniPortalAddr:** All Portal Systems have to have the same value for aDrniPortalAddr.

- 1 j) **aDrniPortalPriority:** All Portal Systems have to have the same value for aDrniPortalPriority.
- 2 k) **aDrniThreePortalSystem:** All Portal Systems have to have the same value for
- 3 aDrniThreePortalSystem.
- 4 l) **aDrniPortalSystemNumber:** All Portal Systems have to have different aDrniPortalSystemNumber
- 5 values, and all of these values have to be in the range 1...3.
- 6 m) **aAggActorAdminKey:** The most significant two bits of the Administrative Key for each
- 7 Aggregator in a Distributed Relay's Aggregator Port is set to the value of
- 8 DRF_Portal_System_Number. The remainder of the bits reflect the physical characteristics of the
- 9 associated Aggregation Ports and they have to be the same for all Portal Systems in the Portal.

10
11 The second class of such information (item b) includes managed objects that control through which Gateway
12 and which Aggregation Ports each Conversation ID passes. For these managed objects, if the information
13 about one Conversation ID is configured differently in different Portal Systems, only that Conversation ID is
14 affected. Therefore, the Portal can operate normally, and the mechanisms that ensure against duplicate
15 delivery or forwarding loops will block the passage of any frames belonging to misconfigured
16 Conversation IDs. In order to detect misconfiguration, so that the blocking is not permanent, the DR
17 Functions can notify the network administrator if the configurations differ. Since these configurations are
18 quite large, a checksum of their contents is exchanged, rather than the configurations, themselves. This
19 method detects differences to a high probability, but not with complete certainty. These managed objects
20 include:

- 21
- 22 n) **aDrniConvAdminGateway[]:** The list that is used to dynamically determine which Gateway
- 23 Conversation ID flows through which Gateway.
- 24 o) **aAggConversationAdminLink[]:** The list that is used to dynamically determine which Port
- 25 Conversation ID flows through which link on the LAG. It is the source of the equivalent mapping
- 26 of Port Conversation IDs to Aggregation Ports used inside the Portal.
- 27

28 DRCP uses its information on which of the expected Portal Systems are either connected or are not
29 connected via IPLs to determine the identity of the emulated Distributed Relay (item c).

30
31 The current operational states of all of the Portal Systems and their Aggregation Ports is exchanged so that
32 each of the DR Functions can determine to which Portal System's Gateway or Aggregator Port each frame is
33 to be delivered (item d). Each DR Function computes a vector specifying exactly which Port
34 Conversation IDs and which Gateway Conversation IDs can pass through each Gateway, Aggregation Port,
35 or IPP. On each IPP, this information is exchanged (item e). Computation of the Gateway Conversation IDs
36 that can pass through each Gateway requires knowledge at each Portal System in a Portal of the operational
37 values of Gateway Vectors on every Gateway in the Portal. In order to avoid sending a 512 Octet vector for
38 each of the Portal System in every DRCPDU exchange, a Gateway Vector database for every Gateway in a
39 Portal is kept by every Portal System. This associates Gateway Vectors with Gateway Sequence numbers,
40 enabling identification of a Gateway Vector in a DRCPDU exchange by its associated Gateway Sequence
41 number. The Gateway Vector database for the Home Portal System should contain at least the two most
42 recent entries. Only when the Gateway Vector changes the new values with its associated Gateway Sequence
43 number would need to be included in a DRCPDU. If there is a difference between two DR Functions'
44 vectors for any given Conversation ID, the output variables are set so that the DR Function will block frames
45 with that Conversation ID. This prevents looping or duplicate delivery of any frame.

46
47 The values of the Keys to be used by the Aggregation Ports on each individual Portal System together with
48 the associated received Operational Keys (DRF_Neighbor_Oper_Partner_Aggregator_Key,
49 DRF_Other_Neighbor_Oper_Partner_Aggregator_Key) are exchanged among the Portal Systems so that a
50 single Aggregator across the Portal is formed (item f). Only the Aggregation Ports that report a
51 Partner_Oper_Key that is equal to the lowest numerical non zero value of the set comprising the values of
52 the DRF_Home_Oper_Partner_Aggregator_Key, the DRF_Neighbor_Oper_Partner_Aggregator_Key and
53 the DRF_Other_Neighbor_Oper_Partner_Aggregator_Key, on each IPP on the Portal System are allowed to
54 be attached to the Portal's Aggregator [item r) in 6.4.14.1].

9.4.1 Establishing the Portal and Distributed Relay

This standard does not specify the creation of Portals automatically. Instead, DRCP compares the network administrator's intentions, as defined by managed objects, to the physical topology of the configured Systems, and if the connected Systems' configurations are compatible, DRCP establishes and enables the operation of the Portal. In order to establish a Distributed Relay across a Portal, a network administrator configures the following managed objects

- a) There may be many Systems in a network, and some or all of the selected Portal Systems may participate in other Portals. Determining which other Portal Systems belong to this Portal System's Portal is accomplished by configuring aDrniPortalAddr (7.4.1.1.4) and aDrniPortalSystemNumber (7.4.1.1.7).
- b) As described in 9.3.2, any point-to-point instance of the MAC Service can be assigned to be an Intra-Portal Link. The particular instances assigned to the use of a DR Function are configured in aDrniIntraPortalLinkList (7.4.1.1.8).
- c) Which Aggregator in each Portal System is to be assigned to this DR Function is configured in aDrniAggregator (7.4.1.1.9).
- d) The methods to be used by the DR Functions to assign frames to Gateway Conversation IDs and Port Conversation IDs are configured in two managed objects, aDrniGatewayAlgorithm (7.4.1.1.13) and aAggPortAlgorithm (7.3.1.1.33).
- e) The prioritized list of assignments of Conversation IDs to Gateways and Aggregation Ports to cover failure modes are configured in several managed objects: aDrniConvAdminGateway[] (7.4.1.1.10), and aAggConversationAdminLink[] (7.3.1.1.35).
- f) If the distributions methods used for the Gateway Conversation IDs and the Port Conversation IDs are the same, the aDrniPortConversationControl (7.4.1.1.23) managed object can be used to assign the Gateway Conversation IDs that are enabled to pass through a DR Function's Gateway based on the Port Conversation IDs assignment towards the DR Function's Aggregator Port, instead of the default operation which is to give the Gateway Conversation ID blocking to the control of the protocols running on the attached network.

9.4.2 DRCPDU transmission, addressing, and protocol identification

Distributed Relay Control Protocol Data Units (DRCPDUs) are transmitted and received using the service provided by an LLC entity that uses, in turn, a single instance of the MAC Service provided at an MAC Service Access Point (MSAP) associated with an IPP. Each DRCPDU is transmitted as a single MAC service request, and received as a single MAC service indication, with the following parameters:

- a) destination address (9.4.2.1)
- b) source address (9.4.2.2)
- c) MSDU
- d) priority (9.4.2.3)

The MSDU of each request and indication comprises a number of octets that provide EtherType protocol identification (9.4.2.4) followed by the DRCPDU proper (9.4.3).

NOTE 1—For the purposes of this standard, the term “LLC entity” includes entities that support protocol discrimination using the EtherType field as specified in IEEE Std 802.

NOTE 2—The complete format of an DRCP frame depends not only on the DRCPDU format, as specified in this clause, but also on the media access method dependent procedures used to support the MAC Service.

9.4.2.1 Destination MAC Address

The destination address for each MAC service request used to transmit a DRCPDU shall be one of the group address specified in Table 9-6 and selected by IPP Managed Objects Class (7.4.2).

Table 9-6—Distributed Relay Control Protocol destination addresses

Assignment	Value
Nearest Customer Bridge group address	01-80-C2-00-00-00
Nearest Bridge group address	01-80-C2-00-00-0E
Nearest non-TPMR Bridge group address	01-80-C2-00-00-03

Its default values shall be the Nearest non-TPMR Bridge group address.

9.4.2.2 Source MAC Address

The source address for each MAC service request used to transmit a DRCPDU shall be an individual address associated with the IPP MSAP at which the request is made.

9.4.2.3 Priority

The priority associated with each MAC Service request should be the default associated with the IPP MSAP.

9.4.2.4 Encapsulation of DRCPDUs in frames

An DRCPDU is encoded in the `mac_service_data_unit` parameter of an `M_UNITDATA.request` or `M_UNITDATA.indication`. The first octets of the `mac_service_data_unit` are a protocol identifier, followed by the DRCPDU, followed by padding octets, if any, as required by the underlying MAC service.

Where the ISS instance used to transmit and receive frames is provided by a media access control method that can support EtherType encoding directly (e.g., is an IEEE 802.3 MAC), the protocol identifier is two octets in length. All DRCPDUs are identified by the DRNI EtherType specified in Table 9-4 and the DRCP subtype specified in Table 9-5.

Where the ISS instance is provided by a media access method that cannot directly support EtherType encoding (e.g., is an IEEE 802.11 MAC), the TPID is encoded according to the rule for a Subnetwork Access Protocol (Clause 10 of IEEE Std 802) that encapsulates Ethernet frames over LLC, and comprises the SNAP header (hexadecimal AA-AA-03) followed by the SNAP PID (hexadecimal 00-00-00) followed by the DRNI protocol's EtherType (hexadecimal 89-52).

9.4.3 DRCPDU structure and encoding

9.4.3.1 Transmission and representation of octets

All DRCPDUs comprise an integral number of octets. The bits in each octet are numbered from 0 to 7, where 0 is the low-order bit. When consecutive octets are used to represent a numerical value, the most significant octet is transmitted first, followed by successively less significant octets.

When the encoding of (an element of) a DRCPDU is depicted in a diagram

- a) Octets are transmitted from top to bottom.

- 1 b) Within an octet, bits are shown with bit 0 to the left and bit 7 to the right, and are transmitted from
2 left to right.
- 3 c) When consecutive octets are used to represent a binary number, the octet transmitted first has the
4 more significant value.
- 5 d) When consecutive octets are used to represent a MAC address, the least significant bit of the first
6 octet is assigned the value of the first bit of the MAC address, the next most significant bit the value
7 of the second bit of the MAC address, and so on through the eighth bit. Similarly, the least
8 significant through most significant bits of the second octet are assigned the value of the ninth
9 through seventeenth bits of the MAC address, and so on for all the octets of the MAC address.

10 **9.4.3.2 DRCPDU structure**

11 The DRCPDU structure shall be as shown in Figure 9-9 and as further described in the following field
12 definitions:

- 13 a) *Subtype*. The Subtype field identifies the specific Protocol being encapsulated. DRCPDUs carry the
14 Subtype value 0x01.
- 15 b) *Version Number*. This identifies the DRCP version; implementations conformant to this version of
16 the standard carry the value 0x01.
- 17 c) **TLV_type = Portal Information**. This field indicates the nature of the information carried in this
18 TLV-tuple. DRNI information is identified by the integer value 1.
- 19 d) *Portal_Information_Length*. This field indicates the length (in octets) of this TLV-tuple, Actor
20 information uses a length value of 16.
- 21 e) *Aggregator_Priority*. The priority assigned to the Actor System ID (by management or
22 administration policy) of the Aggregator attached to the DR Function, encoded as an unsigned
23 integer from aAggActorSystemPriority (7.3.1.1.5).
- 24 f) *Aggregator_ID*. The MAC address component of Actor System ID of the Aggregator attached to the
25 DR Function from aAggActorSystemID (7.3.1.1.4).
- 26 g) *Portal_Priority*. The priority assigned to the Portal's System ID (by management or administration
27 policy), encoded as an unsigned integer from aDrniPortalPriority (7.4.1.1.5).
- 28 h) *Portal_Addr*. The MAC address component of Portal's System ID from aDrniPortalAddr (7.4.1.1.4).
- 29 i) **TLV_type = Portal Configuration Information**. This field indicates the nature of the information
30 carried in this TLV-tuple. *Portal Configuration Information* is identified by the integer value 2.
- 31 j) *Portal_Configuration_Information_Length*. This field indicates the length (in octets) of this TLV-
32 tuple, *Portal Configuration Information* uses a length value of 43.
- 33 k) *Topology_State*. This DR Function's topology related variables for the IPP, encoded as individual
34 bits within a single octet, as follows and as illustrated in Figure 9-10:
- 35 1) *Portal_System_Number* is encoded in bits 0 and 1. The Portal System Number of this DR
36 Function from aDrniPortalSystemNumber (7.4.1.1.7).
- 37 2) *Neighbor_Conf_Portal_System_Number* is encoded in bits 2 and 3. The configured Portal
38 System Number of the Portal System that is attached to this IPP (7.4.1.1.8).
- 39 3) *3_System_Portal* is encoded in bit 4. This bit indicates if this is a Portal System that is part of a
40 Portal consisting of three Portal Systems. It is encoded as 1 for a Portal of three Portal Systems
41 and as 0 otherwise. It is always set equal to aDrniThreePortalSystem (7.4.1.1.6).
- 42 4) *Common_Methods* is encoded in bit 5. It is encoded as 1 when aDrniPortConversationControl
43 (7.4.1.1.23) is set to TRUE and as 0 otherwise.
- 44 5) Bit 6 is reserved for future use. It is set to 0 on transmit and they are ignored on receipt.
- 45 6) *Other_Non_Neighbor* (ONN, 9.4.9) is encoded in bit 7. TRUE (encoded as 1) indicates that the
46 Other Ports Information TLV is not associated with an immediate Neighbor of this Portal
47 System. FALSE (encoded as 0) indicates that the Other Ports Information TLV is an immediate
48 Neighbor on the other IPP on this Portal System. This bit is only used if the Portal Topology
49 contains three Portal Systems. If not, it is transmitted as 0 and ignored on receipt.
- 50
- 51
- 52
- 53
- 54


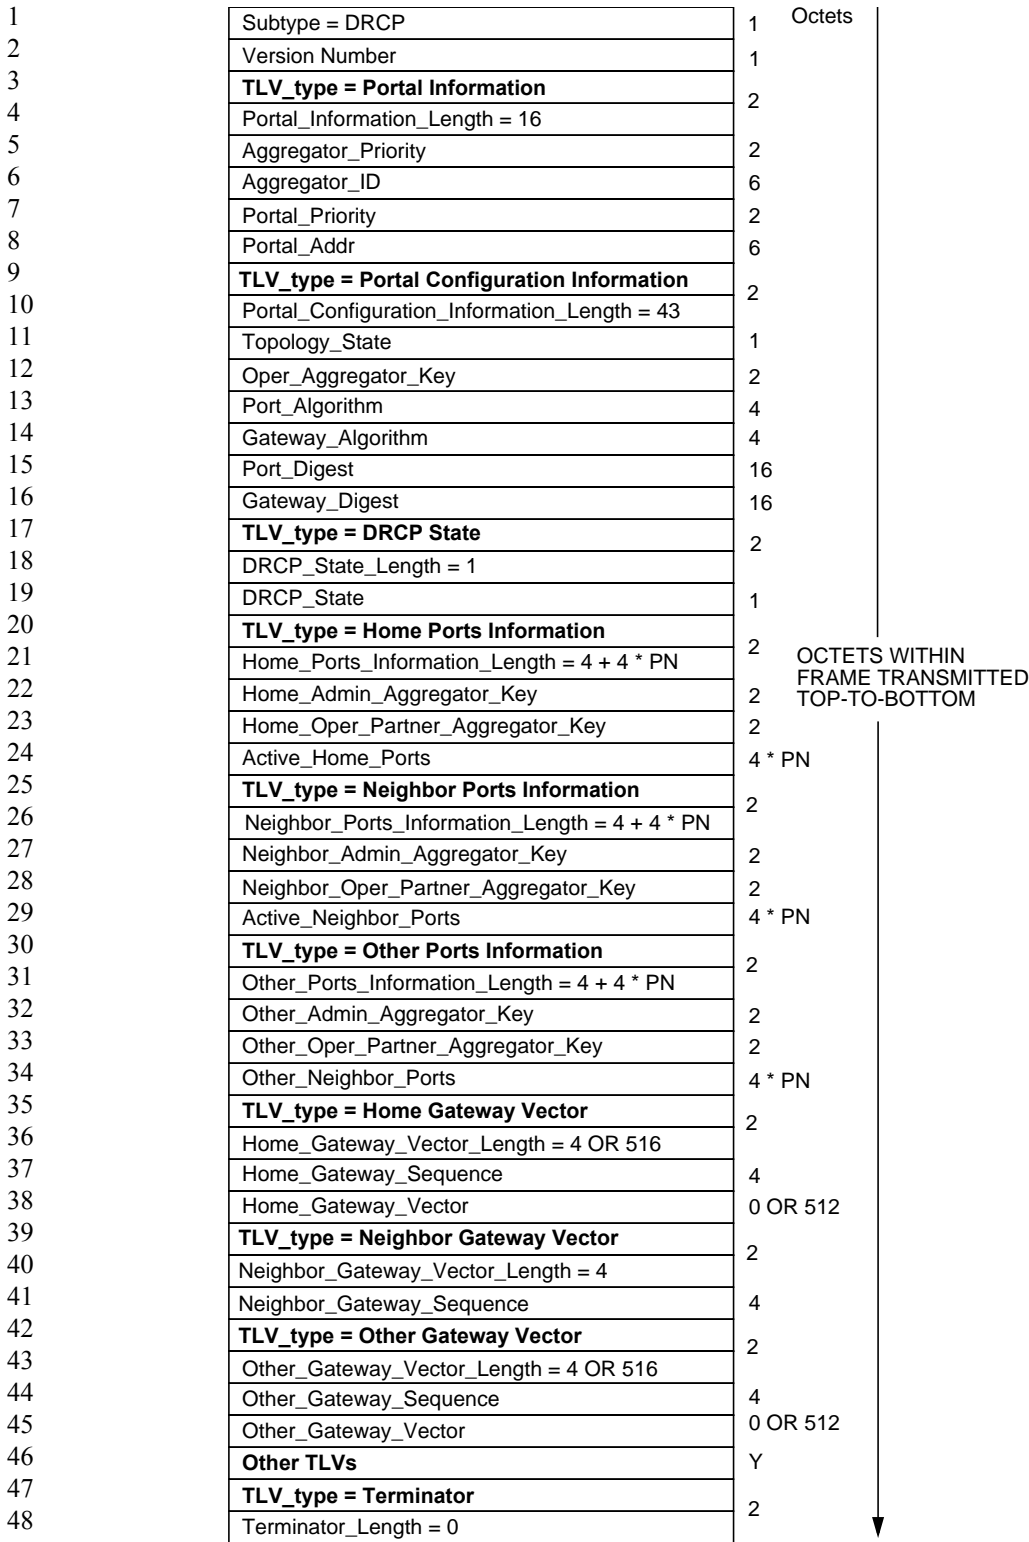


Figure 9-9—DRCPDU structure

BIT							
0	1	2	3	4	5	6	7
Portal_System_Number	Neighbor_Conf_Portals_System_Number	3_System_Portals	Common_Methods	Reserved	ONN		

NOTE 1—Bit ordering within this field is as specified in 9.4.3.1.

Figure 9-10—Bit encoding of the Topology_State field

- l) *Oper_Aggregator_Key*. The current operational Aggregator Key value of the Aggregator attached to the DR Function.
- m) *Port_Algorithm*. The Port algorithm used by this DR Function and Aggregator from aAggPortAlgorithm (7.3.1.1.33). Table 6-4 provides the IEEE 802.1 OUI (00-80-C2) Port Algorithm encodings.
- n) *Gateway_Algorithm*. The Gateway algorithm used by this DR Function from aDrniGatewayAlgorithm (7.4.1.1.13). Table 9-7 provides the IEEE 802.1 OUI (00-80-C2) Gateway Algorithm encodings.

Table 9-7—IEEE Gateway Algorithms

Gateway Algorithm Field	Value
Reserved for IEEE802.1	0
Distribution based on C-VIDs	1
Distribution based on S-VIDs	2
Distribution based on I-SIDs	3
Distribution based on TE-SIDs	4
Distribution based on ECMP Flow Hash	5
Reserved	6-255

- o) *Port_Digest*. The DRF_Home_Conversation_PortList_Digest of this DR Function’s prioritized Port Conversation ID-to-Link Number ID assignments from aAggConversationAdminLink[] (7.3.1.1.35).
- p) *Gateway_Digest*. The DRF_Home_Conversation_GatewayList_Digest of this DR Function’s prioritized Gateway Conversation ID- to-Gateway assignments from aDrniConvAdminGateway[] (7.4.1.1.10).
- q) **TLV_type = DRCP State**. This field indicates the nature of the information carried in this TLV-tuple. *DRCP State* is identified by the integer value 6.
- r) *DRCP_State_Length*. This field indicates the length (in octets) of this TLV-tuple, *DRCP State* uses a length value of 1.
- s) *DRCP_State*. This DR Function’s DRCP variables for the IPP, encoded as individual bits within a single octet, as follows and as illustrated in Figure 9-11:
 - 1) *Home_Gateway* is encoded in bit 0. This flag indicates the operational state of this DR Function’s Gateway. TRUE indicates operational and is encoded as a 1; Non operational is encoded as a 0.
 - 2) *Neighbor_Gateway* is encoded in bit 1. This flag indicates the operational state of the Neighbor DR Function’s Gateway. TRUE indicates operational and is encoded as a 1; Non operational is encoded as a 0.
 - 3) *Other_Gateway* is encoded in bit 2. This flag indicates the operational state of a potential Other neighbor DR Function’s Gateway. TRUE indicates operational and is encoded as a 1;

- 1 Otherwise the flag is encoded as a 0. This bit is only used if the Portal Topology contains three
2 Portal Systems. If not, it is transmitted as 0 and ignored on receipt.
- 3 4) *IPP_Activity* is encoded in bit 3. This flag indicates the Neighbor Portal System's DRCP
4 Activity on this IPP. Active DRCP Neighbor is encoded as a 1; No DRCP activity is encoded as
5 0.
- 6 5) *DRCP_Timeout* is encoded in bit 4. This flag indicates the Timeout control value with regard to
7 this link. Short Timeout is encoded as a 1; Long Timeout is encoded as a 0.
- 8 6) *Gateway_Sync* is encoded in bit 5. If TRUE (encoded as a 1), this DR Function considers this
9 IPP's Neighbor Portal Systems to have their Gateways *IN_SYNC*; i.e., this Portal System's
10 operational vector listing which Portal System's Gateway (if any) is passing each Gateway
11 Conversation ID is in agreement with this IPP's Neighbor Portal Systems' operational vector. If
12 FALSE (encoded as a 0), then this IPP is currently *OUT_OF_SYNC*; i.e., this IPP's Neighbor
13 Portal Systems' operational vector listing which Portal System's Gateway (if any) is passing
14 each Gateway Conversation ID is in disagreement.
- 15 7) *Port_Sync* is encoded in bit 6. If TRUE (encoded as a 1), this DR Function considers this IPP's
16 Neighbor Portal Systems to have their Aggregator Ports *IN_SYNC*; i.e., this Portal System's
17 operational vector listing which Portal System's Aggregation Port (if any) is passing each Port
18 Conversation ID is in agreement with this IPP's Neighbor Portal Systems' operational vector. If
19 FALSE (encoded as a 0), then this IPP is currently *OUT_OF_SYNC*; i.e., this IPP's Neighbor
20 Portal Systems' operational vector listing which Portal System's Aggregation Port (if any) is
21 passing each Port Conversation ID is in disagreement.
- 22 8) *Expired* is encoded in bit 7. If TRUE (encoded as a 1), this flag indicates that the DR Function's
23 DRCPDU Receive machine is in the EXPIRED or DEFAULTED state; if FALSE (encoded as a
24 0), this flag indicates that the DR Function's DRCPDU Receive machine is neither in the
25 EXPIRED nor DEFAULTED state.

26
27 NOTE 2—The received value of Expired state is not used by DRCP; however, knowing its value can be useful when
28 diagnosing protocol problems.

29
30 BIT

0	1	2	3	4	5	6	7
Home_Gateway	Neighbor_Gateway	Other_Gateway	IPP_Activity	DRCP_Timeout	Gateway_Sync	Port_Sync	Expired

31
32
33

34 NOTE 3—Bit ordering within this field is as specified in 9.4.3.1.

35
36 **Figure 9-11—Bit encoding of the DRCP_State field**

- 37
38 t) ***TLV_type = Home Ports Information.*** This field indicates the nature of the information carried in
39 this TLV-tuple. Home Ports information is identified by the integer value 4.
- 40 u) ***Home_Ports_Information_Length.*** This field indicates the length (in octets) of this TLV-tuple.
41 Home Ports information uses a length value of 4 plus 4 times the number of this Portal System's
42 Aggregation Ports that are included in this TLV.
- 43 v) ***Home_Admin_Aggregator_Key.*** The administrative Aggregator Key value of the Aggregator
44 attached to this DR Function from aAggActorAdminKey (7.3.1.1.7).
- 45 w) ***Home_Oper_Partner_Aggregator_Key.*** The operational Partner Aggregator Key associated with
46 this Portal System's Aggregator LAG ID.
- 47 x) ***Active_Home_Ports.*** The list of active Aggregation Ports in this Portal System, provided by the
48 *Drni_Portal_System_State[]* variable, in increasing Port ID (6.3.4) order. The list is controlled by the
49 operation of LACP (listing all the Ports on this Portal System for which LACP is declaring
50 *Actor_Oper_Port_State.Distributing = TRUE*).
- 51 y) ***TLV_type = Neighbor Ports Information.*** This field indicates the nature of the information carried
52 in this TLV-tuple. Neighbor Ports information is identified by the integer value 5.
- 53
54

- 1 z) *Neighbor_Ports_Information_Length*. This field indicates the length (in octets) of this TLV-tuple.
2 Neighbor Ports information uses a length value of 4 plus 4 times the number of the Neighbor
3 Aggregation Ports that are included in this TLV.
- 4 aa) *Neighbor_Admin_Aggregator_Key*. The administrative Aggregator Key value of the Aggregator
5 attached to the Neighbor Portal System.
- 6 ab) *Neighbor_Oper_Partner_Aggregator_Key*. The operational Partner Aggregator Key associated with
7 the Neighbor Portal System's Aggregator LAG ID.
- 8 ac) *Active_Neighbor_Ports*. The list of active Aggregation Ports in the immediate Neighbor Portal
9 System on this IPP, provided by the *Drni_Portal_System_State[]* variable, in increasing Port ID
10 (6.3.4) order. The list is controlled by the operation of LACP (listing all the Ports on the immediate
11 Neighbor Portal System for which LACP is declaring *Actor_Oper_Port_State.Distributing =*
12 *TRUE*).
- 13 ad) ***TLV_type = Other Ports Information***. This field indicates the nature of the information carried in
14 this TLV-tuple. The Other Ports information is identified by the integer value 6. This TLV is only
15 used if the Portal Topology contains three Portal Systems.
- 16 ae) *Other_Ports_Information_Length*. This field indicates the length (in octets) of this TLV-tuple. Other
17 Ports information uses a length value of 4 plus 4 times the number of the Other Portal System's
18 Aggregation Ports that are included in this TLV.
- 19 af) *Other_Admin_Aggregator_Key*. The administrative Aggregator Key value of the Aggregator
20 attached to the Other neighbor Portal System.
- 21 ag) *Other_Oper_Partner_Aggregator_Key*. The operational Partner Aggregator Key associated with the
22 Other neighbor Portal System's Aggregator LAG ID.
- 23 ah) *Other_Neighbor_Ports*. The list of active Aggregation Ports in the Other neighbor Portal System
24 associated with this IPP, provided by the *Drni_Portal_System_State[]* variable, in increasing Port ID
25 (6.3.4) order. The list is controlled by the operation of LACP (listing all the Ports on an optional
26 other Portal System for which LACP is declaring *Actor_Oper_Port_State.Distributing = TRUE*).
- 27 ai) ***TLV_type = Home Gateway Vector***. This field indicates the nature of the information carried in this
28 TLV-tuple. Home Gateway Vector is identified by the integer value 7.
- 29 aj) *Home_Gateway_Vector_Length*. This field indicates the length (in octets) of this TLV-tuple. Home
30 Gateway Vector uses a length value of 4 or 516, the latter when the Home Gateway Vector field is
31 included in the TLV.
- 32 ak) *Home_Gateway_Sequence*. The *DRF_Home_Gateway_Sequence* (9.4.8) sequence number
33 associated with the operational Home Gateway Vector. This corresponds to the first entry in the
34 Home Gateway database.
- 35 al) *Home_Gateway_Vector*. This field contains the value of the
36 *DRF_Home_Gateway_Conversation_Mask* (9.4.8) Boolean vector indexed by increasing Gateway
37 Conversation ID order starting from Gateway Conversation ID 0 up to Gateway Conversation ID
38 4095.
- 39 am) ***TLV_type = Neighbor Gateway Vector***. This field indicates the nature of the information carried in
40 this TLV-tuple. Neighbor Gateway Vector is identified by the integer value 8.
- 41 an) *Neighbor_Gateway_Vector_Length*. This field indicates the length (in octets) of this TLV-tuple.
42 Neighbor Gateway Vector uses a length value of 4.
- 43 ao) *Neighbor_Gateway_Sequence*. The *DRF_Neighbor_Gateway_Sequence* (9.4.9) sequence number
44 associated with the operational Neighbor Gateway Vector.
- 45 ap) ***TLV_type = Other Gateway Vector***. This field indicates the nature of the information carried in this
46 TLV-tuple. Other Gateway Vector is identified by the integer value 9. This TLV is only used if the
47 Portal Topology contains three Portal Systems.
- 48 aq) *Other_Gateway_Vector_Length*. This field indicates the length (in octets) of this TLV-tuple. Other
49 Gateway Vector uses a length value of 4 or 516, the latter when the Other Gateway Vector field is
50 included in the TLV.
- 51 ar) *Other_Gateway_Sequence*. The *DRF_Other_Neighbor_Gateway_Sequence* (9.4.9) sequence
52 number associated with the operational Home Gateway Vector.
- 53 as) *Other_Gateway_Vector*. This field contains the value of the
54 *DRF_Other_Neighbor_Gateway_Conversation_Mask* (9.4.9) Boolean vector indexed by increasing

Gateway Conversation ID order starting from Gateway Conversation ID 0 up to Gateway Conversation ID 4095.

- at) **Other TLVs.** This field contains additional optional TLVs (e.g. 9.4.3.4, 9.4.3.5).
- au) **TLV_type = Terminator.** This field indicates the nature of the information carried in this TLV-tuple. Terminator (end of message) information is identified by the integer value 0.
- av) **Terminator_Length.** This field indicates the length (in octets) of this TLV-tuple. Terminator information uses a length value of 0 (0x00).

NOTE 4—The use of a Terminator_Length of 0 is intentional. In TLV encoding schemes it is common practice for the terminator encoding to be 0 both for the type and the length.

NOTE 5—The Version 1 implementation is guaranteed to be able to receive version N PDUs successfully, although version N PDUs may contain additional information that cannot be interpreted (and will be ignored) by the Version 1 implementation. A crucial factor in ensuring backwards compatibility is that any future version of the protocol is required not to redefine the structure or semantics of information defined for the previous version; it may only add new information elements to the previous set. Hence, in a version N PDU, a Version 1 implementation can expect to find the Version 1 information in exactly the same places as in a Version 1 PDU, and can expect to interpret that information as defined for Version 1.

NOTE 6—The DRCPDU grows in size with the number of Aggregation Ports. A maximum of 90 Aggregation Ports spread across the Portal’s Portal Systems are supported. The minimum number of Aggregation Ports that need to be supported by a Portal is two.

The basic TLV format, applicable to all DRCPDU TLVs, is shown in Figure 9-12.




Figure 9-12—Basic TLV format

The TLV type field occupies the six least significant bits of the first octet of the TLV format. The two most significant bits in the first octet of the TLV format is the two least significant bits of the TLV length field.

NOTE 7—The DRCP TLV Length fields provide the length of the TLV Value fields in contrast to the LACP TLV Length fields which provide the length of the total TLV (including the 2 octets for the TLV Type and TLV Length field). This enables the use of TLVs carrying a TLV Value field of a length up to 1023₄ octets.

Table 9-8 provides a list of the TLVs that are applicable for DRCP.

Table 9-8—Type field values of DRCP TLVs

Type Field	TLV	Support
0	Terminator TLV	Mandatory
1	Portal Information TLV	Mandatory
2	Portal Configuration Information TLV	Mandatory

Table 9-8—Type field values of DRCP TLVs

Type Field	TLV	Support
3	DRCP State TLV	Mandatory
4	Home Ports Information TLV	Mandatory
5	Neighbor Ports Information TLV	Mandatory
6	Other Ports TLV	Mandatory only for Portals of 3 Portal Systems
7	Home Gateway Vector TLV	Mandatory
8	Neighbor Gateway Vector TLV	Mandatory
9	Other Gateway Vector TLV	Mandatory only for Portals of 3 Portal Systems
10	2P Gateway Conversation Vector TLV	Mandatory (see 9.4.3.3 for details)
<u>11</u>	<u>3P Gateway Conversation Vector-1 TLV</u>	<u>Mandatory (see 9.4.3.3 for details)</u>
<u>12</u>	<u>3P Gateway Conversation Vector-2 TLV</u>	<u>Mandatory (see 9.4.3.3 for details)</u>
13	2P Port Conversation Vector TLV	Mandatory (see 9.4.3.3 for details)
<u>14</u>	<u>3P Port Conversation Vector-1 TLV</u>	<u>Mandatory (see 9.4.3.3 for details)</u>
<u>15</u>	<u>3P Port Conversation Vector-2 TLV</u>	<u>Mandatory (see 9.4.3.3 for details)</u>
16	Network/IPL Sharing Method TLV	Optional (see 9.4.3.4 for details)
17	Network/IPL Sharing Encapsulation TLV	Optional (see 9.4.3.4 for details)
18	Organization Specific TLV	Optional (see 9.4.3.5 for details)
19 5-63	Reserved for IEEE 802.1	

9.4.3.3 Conversation Vector TLVs

Table 9-9 provides the list of the Conversation Vector TLVs. These TLVs shall be present in a DRCPDU

Table 9-9—Type field values of Conversation Vector TLVs

TLV	Type Field
2P Gateway Conversation Vector TLV	10
3P Gateway Conversation Vector-1 TLV	11
3P Gateway Conversation Vector-2 TLV	12
2P Port Conversation Vector TLV	13
3P Port Conversation Vector-1 TLV	14
3P Port Conversation Vector-2 TLV	15

only when certain conditions are met. In particular the 2P Gateway Conversation Vector TLV shall only be present when GatewayConversationTransmit == TRUE and aDrniThreePortalSystem == FALSE, the 3P Gateway Conversation Vector-1 TLV accompanied by the 3P Gateway Conversation Vector-2 TLV shall only be present when GatewayConversationTransmit == TRUE and aDrniThreePortalSystem == TRUE, and the 2P Port Conversation Vector TLV shall only be present when PortConversationTransmit == TRUE and aDrniThreePortalSystem == FALSE, and the 3P Port Conversation Vector-1 TLV accompanied by the 3P Port Conversation Vector-2 TLV shall only be present when PortConversationTransmit == TRUE and aDrniThreePortalSystem == TRUE. See the DRCPDU Transmit machine (9.4.19) for details regarding the presence of the Conversation Vector TLVs in a DRCPDU.

9.4.3.3.1 2P Gateway Conversation Vector TLV

The 2P Gateway Conversation Vector TLV structure shall be as shown in Figure 9-13 and as further described in the following field definitions:

TLV_type = 2P Gateway Conversation Vector	2
2P_Gateway_Conversation_Vector_Length = 512	
Gateway_Conversation 2P Gateway Conversation Vect	512

Figure 9-13—2P Gateway Conversation Vector TLV

- a) TLV_type = 2P Gateway Conversation Vector. This field indicates the nature of the information carried in this TLV-tuple. 2P Gateway Conversation Vector is identified by the integer value 10.
- b) 2P_Gateway_Conversation_Vector_Length. This field indicates the length (in octets) of this TLV-tuple. 2P Gateway Conversation Vector uses a length value of 512 or 1024.
- c) Gateway_Conversation 2P Gateway Conversation Vector. The contents of this field depend on the value of the aDrniThreePortalSystem flag. If aDrniThreePortalSystem == FALSE, the field provides the operational Drni_Portal_System_Gateway_Conversation (9.3.4.2) Boolean vector. Drni_Portal_System_Gateway_Conversation 2P Gateway Conversation Vector is encoded as a 512-octet Boolean vector, indexed by Gateway Conversation ID, based on the Drni_Portal_System_Gateway_Conversation. The first bit indicates if Gateway Conversation ID 0 is allowed to pass through this DR Function’s Gateway, the second bit indicates if Gateway Conversation ID 1 is allowed to pass through this DR Function’s Gateway, and so on until the final

bit which indicates if Gateway Conversation ID 4095 is allowed to pass through this DR Function's Gateway. ~~If aDmriThreePortalSystem == TRUE, the field provides the operational Dmri_Gateway_Conversation (9.4.7) vector listing which Portal System's Gateway (if any) is passing each Gateway Conversation ID. Dmri_Gateway_Conversation is encoded as a 1024-octet vector as sequence of Portal System Numbers (0 for none), indexed by Gateway Conversation ID. The first two bits provide the Portal System Number of the Portal System that is passing Gateway Conversation ID 0, the second two bits provide the Portal System Number of the Portal System that is passing Gateway Conversation ID 1, and so on until the final two bits which provide the Portal System Number of the Portal System that is passing Gateway Conversation ID 4095.~~

9.4.3.3.2 3P Gateway Conversation Vector-1 TLV

The 3P Gateway Conversation Vector-1 TLV structure shall be as shown in Figure 9-14 and as further described in the following field definitions:

TLV_type = 3P Gateway Conversation Vector-1	2
3P_Gateway_Conversation_Vector_1_Length = 512	
Gateway_Conversation 3P Gateway Conversation Vecto	512

Figure 9-14—3P Gateway Conversation Vector-1 TLV

- TLV_type = 3P Gateway Conversation Vector-1.* This field indicates the nature of the information carried in this TLV-tuple. *3P Gateway Conversation Vector-1* is identified by the integer value 11.
- 3P_Gateway_Conversation_Vector_1_Length.* This field indicates the length (in octets) of this TLV-tuple. *3P Gateway Conversation Vector-1* uses a length value of 512.
- 3P_Gateway_Conversation_Vector-1.* The field provides the operational *Dmri_Gateway_Conversation (9.4.7)* vector listing which Portal System's Gateway (if any) is passing each of the first 2048 Gateway Conversation IDs (Gateway Conversation ID 0 to Gateway Conversation ID 2047). *3P_Gateway_Conversation_Vector-1* is encoded as a 512-octet vector as sequence of Portal System Numbers (0 for none), indexed by Gateway Conversation ID, based on the *Dmri_Gateway_Conversation*. The first two bits provide the Portal System Number of the Portal System that is passing Gateway Conversation ID 0, the second two bits provide the Portal System Number of the Portal System that is passing Gateway Conversation ID 1, and so on until the final two bits which provide the Portal System Number of the Portal System that is passing Gateway Conversation ID 2047.

9.4.3.3.3 3P Gateway Conversation Vector-2 TLV

The 3P Gateway Conversation Vector-2 TLV structure shall be as shown in Figure 9-15 and as further described in the following field definitions:

TLV_type = 3P Gateway Conversation Vector-2	2
3P_Gateway_Conversation_Vector_2_Length = 512	
3P Gateway Conversation Vector-2	512

Figure 9-15—3P Gateway Conversation Vector-2 TLV

- TLV_type = 3P Gateway Conversation Vector-2.* This field indicates the nature of the information carried in this TLV-tuple. *3P Gateway Conversation Vector-2* is identified by the integer value 12.

- b) *3P_Gateway_Conversation_Vector_2_Length*. This field indicates the length (in octets) of this TLV-tuple. *3P_Gateway_Conversation_Vector-2* uses a length value of 512.
- c) *3P_Gateway_Conversation_Vector-2*. The field provides the operational *Drni_Gateway_Conversation* (9.4.7) vector listing which Portal System’s Gateway (if any) is passing each of the last 2048 Gateway Conversation IDs (Gateway Conversation ID 2048 to Gateway Conversation ID 4095). *3P_Gateway_Conversation_Vector-2* is encoded as a 512-octet vector as sequence of Portal System Numbers (0 for none), indexed by Gateway Conversation ID, based on the *Drni_Gateway_Conversation*. The first two bits provide the Portal System Number of the Portal System that is passing Gateway Conversation ID 2048, the second two bits provide the Portal System Number of the Portal System that is passing Gateway Conversation ID 2049, and so on until the final two bits which provide the Portal System Number of the Portal System that is passing Gateway Conversation ID 4095.

9.4.3.3.4 2P Port Conversation Vector TLV

The 2P Port Conversation Vector TLV structure shall be as shown in Figure 9-16 and as further described in the following field definitions:

TLV_type = 2P Port Conversation Vector	2
2P_Port_Conversation_Vector_Length = 512	
Port_Conversation 2P Port Conversation Vector	512

Figure 9-16—2P Port Conversation Vector TLV

- a) *TLV_type = 2P Port Conversation Vector*. This field indicates the nature of the information carried in this TLV-tuple. *2P Port Conversation Vector* is identified by the integer value 1~~3~~.
- b) *2P_Port_Conversation_Vector_Length*. This field indicates the length (in octets) of this TLV-tuple. *2P_Port_Conversation_Vector* uses a length value of 512 ~~or 1024~~.
- c) *2P_Port_Conversation_Vector*. ~~The contents of this field depend on the value of the aDrniThreePortalSystem flag. If aDrniThreePortalSystem == FALSE, the field provides the operational Drni_Portal_System_Port_Conversation (9.3.4.2) Boolean vector. 2P_Port_Conversation_Vector Drni_Portal_System_Port_Conversation is encoded as a 512-octet Boolean vector, indexed by Port Conversation ID, based on the Drni_Portal_System_Port_Conversation. The first bit indicates if Port Conversation ID 0 is allowed to pass through this DR Function’s Aggregator, the second bit indicates if Port Conversation ID 1 is allowed to pass through this DR Function’s Aggregator, and so on until the final bit which indicates if Port Conversation ID 4095 is allowed to pass through this DR Function’s Aggregator. If aDrniThreePortalSystem == TRUE, the field provides the operational Drni_Port_Conversation (9.4.7) vector listing which Portal System’s Aggregator (if any) is passing each Port Conversation ID. It is encoded in a 1024-octet vector as sequence of Portal System Numbers (0 for none), indexed by Port Conversation ID. The first two bits provide the Portal System Number of the Portal System that is passing Port Conversation ID 0, the second two bits provide the Portal System Number of the Portal System that is passing Port Conversation ID 1, and so on until the final two bits which provide the Portal System Number of the Portal System that is passing Port Conversation ID 4095.~~

9.4.3.3.5 3P Port Conversation Vector-1 TLV

The 3P Port Conversation Vector-1 TLV structure shall be as shown in Figure 9-17 and as further described in the following field definitions:

TLV_type = 3P Port Conversation Vector-1	2
3P_Port_Conversation_Vector_1_Length = 512	
Port_Conversation 3P Port Conversation Vector-1	512

Figure 9-17—3P Port Conversation Vector-1 TLV

- TLV_type = 3P Port Conversation Vector-1.* This field indicates the nature of the information carried in this TLV-tuple. *3P Port Conversation Vector-1* is identified by the integer value 14.
- 3P Port Conversation Vector 1 Length.* This field indicates the length (in octets) of this TLV-tuple. *3P Port Conversation Vector-1* uses a length value of 512.
- 3P Port Conversation Vector-1.* The field provides the operational *Drni Port Conversation* (9.4.7) vector listing which Portal System's Aggregator (if any) is passing each of the first 2048 Port Conversation IDs (Port Conversation ID 0 to Port Conversation ID 2047). *3P Port Conversation Vector-1* is encoded in a 512-octet vector as sequence of Portal System Numbers (0 for none), indexed by Port Conversation ID, based on the *Drni Port Conversation*. The first two bits provide the Portal System Number of the Portal System that is passing Port Conversation ID 0, the second two bits provide the Portal System Number of the Portal System that is passing Port Conversation ID 1, and so on until the final two bits which provide the Portal System Number of the Portal System that is passing Port Conversation ID 2047.

9.4.3.3.6 3P Port Conversation Vector-2 TLV

The 3P Port Conversation Vector-2 TLV structure shall be as shown in Figure 9-18 and as further described in the following field definitions:

TLV_type = 3P Port Conversation Vector-2	2
3P_Port_Conversation_Vector_2_Length = 512	
Port_Conversation 3P Port Conversation Vector-2	512

Figure 9-18—3P Port Conversation Vector-2 TLV

- TLV_type = 3P Port Conversation Vector-2.* This field indicates the nature of the information carried in this TLV-tuple. *3P Port Conversation Vector-2* is identified by the integer value 15.
- 3P Port Conversation Vector 2 Length.* This field indicates the length (in octets) of this TLV-tuple. *3P Port Conversation Vector-2* uses a length value of 512.
- 3P Port Conversation Vector-2.* The field provides the operational *Drni Port Conversation* (9.4.7) vector listing which Portal System's Aggregator (if any) is passing each of the last 2048 Port Conversation IDs (Port Conversation ID 2048 to Port Conversation ID 4095). *3P Port Conversation Vector-2* is encoded in a 512-octet vector as sequence of Portal System Numbers (0 for none), indexed by Port Conversation ID, based on the *Drni Port Conversation*. The first two bits provide the Portal System Number of the Portal System that is passing Port Conversation ID 2048, the second two bits provide the Portal System Number of the Portal System that is passing Port Conversation ID 2049, and so on until the final two bits which provide the Portal System Number of the Portal System that is passing Port Conversation ID 4095.

9.4.3.4 Network/IPL sharing TLVs

These TLVs are only required when the Network/IPL sharing method used is one of Network / IPL sharing by tag (9.3.2.2), or Network / IPL sharing by encapsulation (9.3.2.3) in order to ensure consistent configuration among the Portal Systems. The Network / IPL sharing by time (9.3.2.1) method requires the exchange of the Network/IPL Sharing Method TLV (9.4.3.4.1) but not the Network/IPL Sharing Encapsulation TLV (9.4.3.4.2).

NOTE—No Network/IPL sharing TLVs are required when the Network/IPL sharing method used is the physical or aggregated method in 9.3.2.

Table 9-10 provides a list of the TLVs that are applicable for Network/IPL sharing methods. They are

Table 9-10—Type field values of Network/IPL sharing TLVs

TLV	Type Field
Network/IPL Sharing Method TLV	16 2
Network/IPL Sharing Encapsulation TLV	17 3

required to support the functionality described in 9.4.20.

9.4.3.4.1 Network/IPL Sharing Method TLV

The Network/IPL Sharing Method TLV structure shall be as shown in Figure 9-19 and as further described in the following field definitions:

TLV_type = Network/IPL Sharing Method	2
Network/IPL_Sharing_Method_Length = 4	
DRF_Home_Network/IPL_Sharing_Method	4

Figure 9-19—Network/IPL Sharing Method TLV

- TLV_type = Network/IPL Sharing Method TLV*. This field indicates the nature of the information carried in this TLV-tuple. The Network/IPL sharing TLVs is identified by the integer value 16~~2~~.
- Network/IPL_Sharing_Method_Length*. This field indicates the length (in octets) of this TLV-tuple. The Network/IPL sharing TLVs uses a length value of 4.
- DRF_Home_Network/IPL_Sharing_Method*. This field contains the value representing the Network/IPL sharing method that is used to transport IPL frames to the Neighbor Portal System on this IPP when the IPL and network link are sharing the same physical link. It consists of the three-octet OUI or Company Identifier (CID) identifying the organization which is responsible for this encapsulation and one following octet used to identify the encapsulation method defined by that organization. Always set equal to aDrniEncapsulationMethod (7.4.1.1.17). A value of 1 indicates that Network / IPL sharing by time (9.3.2.1) is used. A value of 2 indicates that the encapsulation method used is the same as the one used by network frames and that Network / IPL sharing by tag

(9.3.2.2) is used. Table 9-11 provides the IEEE 802.1 OUI (00-80-C2) encapsulation method encodings.

Table 9-11—IEEE encapsulation methods

Encapsulation Method Field	Value
IPL is using a separate physical or Aggregation link	0
Network / IPL sharing by time (9.3.2.1)	1
Network / IPL sharing by tag (9.3.2.2)	2
IEEE802.1Q I-TAG based encapsulation	3
IEEE802.1Q B-VLAN based encapsulation	4
IETF Pseudowire based encapsulation	5
Reserved	6-255

9.4.3.4.2 Network/IPL Sharing Encapsulation TLV

The Network/IPL Sharing Encapsulation TLV structure shall be as shown in Figure 9-20 and as further described in the following field definitions:

TLV_type = Network/IPL Sharing Encapsulation	2
Network/IPL_Sharing_Encapsulation_Length = 32	
DRF_Home_Network/IPL_IPLEncap_Digest	16
DRF_Home_Network/IPL_NetEncap_Digest	16

Figure 9-20—Network/IPL Sharing Encapsulation TLV

- TLV_type = Network/IPL Sharing Encapsulation TLV.* This field indicates the nature of the information carried in this TLV-tuple. The Network/IPL sharing TLVs is identified by the integer value 173.
- Network/IPL_Sharing_Encapsulation_Length.* This field indicates the length (in octets) of this TLV-tuple. The Network/IPL sharing TLVs uses a length value of 32.
- DRF_Home_Network/IPL_IPLEncap_Digest.* This field contains the value of the MD5 digest computed from aDrniIPLEncapMap (7.4.1.1.18) for exchange with the Neighbor Portal System on the IPL.
- DRF_Home_Network/IPL_NetEncap_Digest.* This field contains the value of the MD5 digest computed from aDrniNetEncapMap (7.4.1.1.19) for exchange on the shared network link.

9.4.3.5 Organization-Specific TLV

These optional Organization-Specific TLVs are provided to allow different organizations, such as IEEE 802.1, ITU-T, IETF, as well as individual software and equipment vendors, to define TLVs that advertise

additional information to Neighbor Portal Systems. The Organization-Specific TLV structure shall be as shown in Figure 9-21 and as further described in the following field definitions:

TLV_type = Organization-Specific	2
Organization_Specific_Length = LL	
OUI/CID	3
Subtype	7
Value (Optional)	LL-10

Figure 9-21—Organization-Specific TLV

- a) *TLV_type = Organization-Specific TLV*. This field indicates the nature of the information carried in this TLV-tuple. The Organization-Specific TLV is identified by the integer value 184.
- b) *Organization_Specific_Length*. This field indicates the length (in octets) of this TLV-tuple. The Organization-Specific TLV uses a length value of LL.
- c) *OUI/CID*. This field contains the 3-byte long Organizationally Unique Identifier or Company Identifier, obtainable from IEEE.
- d) *Subtype*. This field contains a Subtype value, so that an additional OUI/CID will not be required if more Organization-Specific TLVs are required by an owner of an OUI/CID.
- e) *Value*. This field contains the information that will be communicated to a Neighbor Portal System.

The following restrictions apply:

- f) Information transmitted in an Organization-Specific TLV shall not require the recipient to violate any requirement in this standard;
- g) Information transmitted in one Organization-Specific TLV shall not be used to provide a means for sending messages that are larger than would fit within a single DRCPDU.

9.4.4 DRCP Control Parser/Multiplexer

There can be up to two DRCP Control Parser/Multiplexers on a Portal System, one for each IPP. Each DRCP Control Parser/Multiplexer N, where N identifies one of the two IPPs, is an instance of the Protocol Parser/Multiplexer described in 6.2.7. Specifically:

- a) The *DownPort* of the Protocol Parser/Multiplexer is *MacIppN* (9.3.4.1) for DRCP Control Parser/Multiplexer N.
- b) The *ControlPort* of the Protocol Parser/Multiplexer is *DRCPCtrlMuxN* (9.3.4.1) for DRCP Control Parser/Multiplexer N.
- c) The *DataPort* of the Protocol Parser/Multiplexer is *IppN* (9.3.4.1) for DRCP Control Parser/Multiplexer N.
- d) The *IsControlFrame* function is defined in 9.4.4.1.

9.4.4.1 Control Parser state diagram

9.4.4.1.1 Control Parser Function

IsControlFrame

Returns Boolean value: (DA == DRCP_Protocol_DA && Length/Type == DRCP_Type && Subtype == DRCP_subtype)

9.4.4.1.2 Constants

DRCP_Type

The value of the Distributed Relay Control Protocol EtherType field.

DRCP_subtype

The value of the Subtype field for the DRCP.

Value: Integer

1

9.4.4.1.3 Variables

DA

SA

mac_service_data_unit

priority

The parameters of the M_UNITDATA.indication primitive.

DRCP_Protocol_DA

One of the addresses selected from Table 9-6 determined by the setting of the IPP Managed Objects Class managed object (7.4.2).

Value: 48 bits.

Length/Type

The value of the Length/Type field in a received frame.

Value: Integer

Subtype

The value of the octet following the Length/Type field.

Value: Integer

9.4.5 DRCP state machine overview

The operation of the protocol is controlled by a number of state machines, each of which performs a distinct function. These state machines are for the most part described on a per-IPP basis; any deviations from per-IPP description are highlighted in the text. Events (such as expiration of a timer or received DRCPDUs) may cause state transitions and also cause actions to be taken; those actions may include the need for transmission of a DRCPDU containing repeated or new information. Periodic and event-driven transmissions are controlled by the state of a Need-To-Transmit (NTTDRCPDU) variable, generated by the state machines as necessary.

The state machines are as follows:

- a) *DRCPDU Receive machine (DRX—9.4.14)*. This state machine receives DRCPDUs from the Neighbor Portal System on this IPP, records the information contained, and times it out using either Short Timeouts or Long Timeouts, according to the setting of DRCP_Timeout. It evaluates the incoming information from the Neighbor Portal System to determine whether the Home and Neighbor have both agreed upon the protocol information exchanged to the extent that the Home Portal System can now be safely used, either in Portal with other Portal Systems or as an individual Portal; if not, it asserts NTTDRCPDU in order to transmit fresh protocol information to the Neighbor Portal System. If the protocol information from the Neighbor Portal System times out, the DRCPDU Receive machine installs default parameter values for use by the other state machines.
- b) *DRCP Periodic Transmission machine (Periodic—9.4.15)*. This state machine determines the period that the Home Portal System and its Neighbors will exchange DRCPDUs in order to maintain the Portal.
- c) *Portal System machine (PS—9.4.16)*. This state machine is responsible to update the operational status of all the Gateways and Aggregation Ports in the Portal based on local information and DRCPDUs received on the Home Portal System's IPPs. This state machine is per Portal System.

- 1 d) *DRNI Gateway and Aggregator machines (DGA—9.4.17)*. These state machines are responsible for
2 configuring the Gateway Conversation IDs which are allowed to pass through this DR Function's
3 Gateway and the Port Conversation IDs which are allowed to be distributed through this DR
4 Function's Aggregator. These state machines are per Portal System.
- 5 e) *DRNI IPP machines (IPP—9.4.18)*. These state machines are responsible for configuring the
6 Gateway Conversation IDs and the Port Conversation IDs which are allowed to pass through this
7 DR Function's IPPs.
- 8 f) *DRCPDU Transmit machine (DTX—9.4.19)*. This state machine handles the transmission of
9 DRCPDUs, both on demand from the other state machines, and on a periodic basis.

10
11 Figure 9-22 illustrates the relationships among these state machines and the flow of information between
12 them. The set of arrows labeled Neighbor State Information represents new Neighbor information, contained
13 in an incoming DRCPDU or supplied by administrative default values, being fed to each state machine by
14 the DRCPDU Receive machine. The set of arrows labeled Home State Information represents the flow of
15 updated Home state information between the state machines. Transmission of DRCPDUs occurs either as a
16 result of the Periodic machine determining the need to transmit a periodic DRCPDU, or as a result of
17 changes to the Home's state information that need to be communicated to the Neighbors. The need to
18 transmit a DRCPDU is signaled to the DRCPDU Transmit machine by asserting NTTDRCPDU. The
19 remaining arrows represent shared variables in the state machine description that allow a state machine to
20 cause events to occur in another state machine.

21 22 **9.4.6 Constants**

23
24 All timers specified in this subclause have an implementation tolerance of ± 250 ms.

25 Drni_Fast_Periodic_Time

26 The number of seconds between periodic transmissions using Short Timeouts.

27 Value: Integer

28 1

29 Drni_Slow_Periodic_Time

30 The number of seconds between periodic transmissions using Long Timeouts.

31 Value: Integer

32 30

33 Drni_Short_Timeout_Time

34 The number of seconds before invalidating received DRCPDU information when using Short
35 Timeouts ($3 \times \text{Drni_Fast_Periodic_Time}$).

36 Value: Integer

37 3

38 Drni_Long_Timeout_Time

39 The number of seconds before invalidating received DRCPDU information when using Long
40 Timeouts ($3 \times \text{Drni_Slow_Periodic_Time}$).

41 Value: Integer

42 90

43
44
45 NOTE—Timing out DRCPDU exchanges is the method of last resort for detecting IPL failures and shifting data flows to
46 other Portal Systems. The MAC_Operational status generated by the link hardware is the first choice for link failure
47 detection. IEEE Std 802.1Q Clause 18 Connectivity Fault Management provides a method for detecting link failures
48 (also indicated via the MAC_Operational status) when hardware means are not applicable.

49 50 **9.4.7 Variables associated with the Distributed Relay**

51 Drni_Aggregator_ID

52 The MAC address component of the System Identifier of the Aggregator associated with this
53 Portal. Always set equal to aAggActorSystemID (7.3.1.1.4). Transmitted in DRCPDUs.

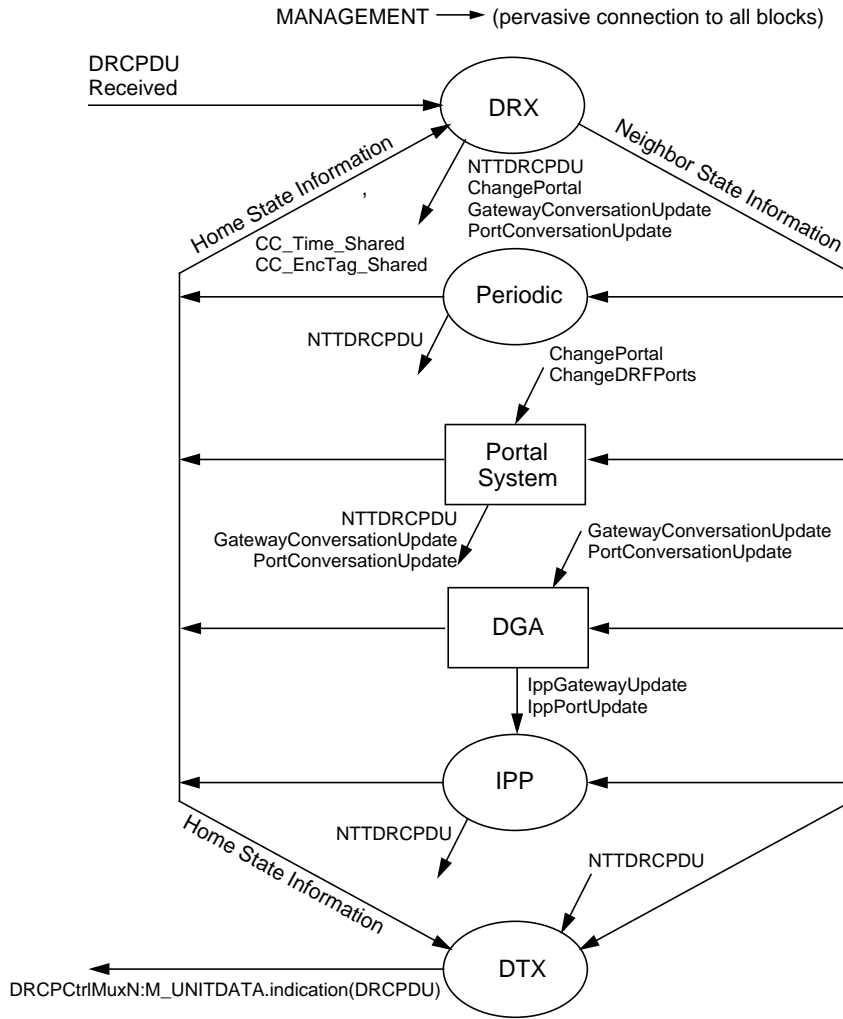


Figure 9-22—Interrelationships among state machines

Value: 48 bits

Drni_Aggregator_Priority

The System Priority of the Aggregator associated to this Portal. Always set equal to aAggActorSystemPriority (7.3.1.1.5). Transmitted in DRCPPDUs.

Value: Integer

Assigned by administrator or System policy.

Drni_Gateway_Conversation

Operational vector listing which Portal System's Gateway (if any) is passing each Gateway Conversation ID.

Value: sequence of Portal System Numbers (0 for none), indexed by Gateway Conversation ID.
Value computed from aDrniConvAdminGateway[] and Drni_Portal_System_State[] upon initialization and whenever the managed object or variable changes.

Drni_Port_Conversation

Operational vector listing which Portal System (if any) is passing each Port Conversation ID.

Value: sequence of Portal System Numbers (0 for none), indexed by Port Conversation ID.
Value computed from Conversation_PortList[] and Drni_Portal_System_State[] upon initialization and whenever the managed object or variable changes.

1 Drni_Portal_Addr
2 The MAC address component of the System Identifier of the Portal. Always set equal to
3 aDrniPortalAddr (7.4.1.1.4). Transmitted in DRCPDUs.
4 Value: 48 bits
5 Assigned by administrator or System policy.
6 Drni_Portal_Priority
7 The System Priority of the Portal. Always set equal to aDrniPortalPriority (7.4.1.1.5).
8 Transmitted in DRCPDUs.
9 Value: Integer
10 Assigned by administrator or System policy.
11 Drni_Three_System_Portal
12 A Boolean indicating if this is a Portal System that is part of a Portal consisting of three Portal
13 Systems. Always set equal to aDrniThreePortalSystem (7.4.1.1.6). Transmitted in DRCPDUs.
14 Value: Boolean
15 Assigned by administrator or System policy.

16 **9.4.8 Per-DR Function variables**

17
18
19 ChangeDRFPorts
20 This variable tracks the operational state of the Gateway, the Gateway Vector, and all
21 Aggregation Ports associated to this Portal System and is set to TRUE when any of them
22 changes. This variable is also set to TRUE if new values for the
23 Drni_Conversation_GatewayList[] or aAggConversationAdminLink[] are initiated.
24 Value: Boolean
25 ChangePortal
26 This variable is set to TRUE when the DRF_Neighbor_Oper_DRCP_State.IPP_Activity on
27 any IPP on this Portal System changes. The variable can also be set to TRUE by the
28 recordPortalConfValues function.
29 Value: Boolean
30 Drni_Common_Methods
31 A flag indicating whether the Gateway and the Port Algorithms use the same methods for
32 frame distribution across the Portal Systems within a Portal. Always set equal to
33 aDrniPortConversationControl (7.4.1.1.23). Transmitted in DRCPDUs.
34 Value: Boolean
35 Drni_Conversation_GatewayList[]
36 An array of 4096 lists, indexed by Gateway Conversation ID, that determines which Gateway
37 in this Portal carries which Gateway Conversation ID. Each item in the array is a list of
38 Gateways in this Portal, in priority order from most desired to least desired, for carrying the
39 indexed Gateway Conversation ID. Assigned by administrator or system policy. Always set
40 equal to aDrniConvAdminGateway[] (7.4.1.1.10).
41 Value: sequence of Port IDs
42 Drni_Portal_System_State[]
43 The states of all Portal Systems in this Portal, indexed by Portal System Number.
44 Value: For each Portal System, a Boolean flag indicating the operational state of the current
45 Portal System's Gateway (TRUE indicates operational), the Gateway Boolean vector of the
46 Portal System, a List (perhaps empty) of the Port IDs of the operational Aggregation Ports in
47 that Portal System, and the identity of the IPP, if any, from which the Portal System's state was
48 obtained.
49 This variable is set by the updatePortalState function (9.4.11). Transmitted in DRCPDUs.
50 DRF_Home_Admin_Aggregator_Key
51 The administrative Aggregator Key value associated with this Portal System's Aggregator.
52 Transmitted in DRCPDUs.
53 Value: Integer
54

Assigned by administrator or System policy. The DRF_Home_Admin_Aggregator_Key is configured and has to be different for each Portal System. Specifically the two most significant bits have to be different in each Portal System. The lower 14 bits may be any value, have to be the same in each Portal System within the same Portal, and have a default of zero.

Assigned by administrator or System policy.

DRF_Home_Conversation_GatewayList_Digest

A digest of aDrniConvAdminGateway[] (7.4.1.1.10), configured in this DR Function, for exchange with the Neighbor Portal Systems. The digest, is a 16-octet MD5 fingerprint [see IETF RFC 1321 (1992)] created from the aDrniConvAdminGateway[]. To calculate the digest, aDrniConvAdminGateway[] is considered to contain 4096 consecutive elements, where each element contains a list of Portal System Numbers, encoded as binary numbers in the order of preference, highest to lowest, followed by the Gateway Conversation ID. The first element of the table contains the prioritized list of Portal System Numbers assigned to Gateway Conversation ID 0, the second element the prioritized list of Portal System Numbers assigned to Gateway Conversation ID 1, the third element the prioritized list of Portal System Numbers assigned to Gateway Conversation ID 2, and so on, with the last element containing the prioritized list of Portal System Numbers assigned to Gateway Conversation ID 4095. This variable is referenced by the DRCPDU.

Value: MD5 digest

DRF_Home_Conversation_PortList_Digest

A digest of aAggConversationAdminLink[] (7.3.1.1.35), configured in this DR Function, for exchange with the Neighbor Portal Systems. Always set equal to Actor_Conversation_LinkList_Digest (6.6.2.1). Transmitted in DRCPDUs.

Value: MD5 digest

DRF_Home_Gateway_Algorithm

The gateway algorithm used by this DR Function to assign frames to Gateway Conversation IDs. Always set equal to the aDrniGatewayAlgorithm (7.4.1.1.13). Transmitted in DRCPDUs.

Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining this algorithm followed by one octet identifying this specific algorithm).

DRF_Home_Gateway_Conversation_Mask

The operational Boolean Gateway vector, indexed by Gateway Conversation ID, indicating whether the indexed Gateway Conversation ID is enabled to pass through the Home Gateway (FALSE = blocked). Its value is updated by the updateDRFHomeState function and is stored as the first entry in the Gateway Vector database for the Home Portal System. Transmitted in DRCPDUs.

Value: sequence of Boolean values, indexed by Gateway Conversation ID.

DRF_Home_Gateway_Sequence

The sequence number of the operational DRF_Home_Gateway_Conversation_Mask vector. It is initialized to zero and is updated by the updateDRFHomeState function. Transmitted in DRCPDUs.

Value: Integer

DRF_Home_Port_Algorithm

The port algorithm used by this DR Function to assign frames to Port Conversation IDs. Always set equal to the associated Aggregator's aAggPortAlgorithm (7.3.1.1.33). Transmitted in DRCPDUs.

Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining this algorithm followed by one octet identifying this specific algorithm).

DRF_Home_Oper_Aggregator_Key

The operational Aggregator Key value associated with this Portal System's Aggregator. Its value is computed by the updateKey function (9.4.11). Transmitted in DRCPDUs.

Value: Integer

DRF_Home_Oper_Partner_Aggregator_Key

1 The operational Partner Aggregator Key associated with this Portal System's Aggregator LAG
2 ID. Transmitted in DRCPDUs.
3 Value: Integer
4 DRF_Home_State
5 The operational state of this DR Function. Transmitted in DRCPDUs.
6 Value: Boolean flag indicating the operational state of this Portal System's Gateway (TRUE
7 indicates operational), the Boolean Gateway Vector associated with the most recent entry in the
8 Gateway Vector database for this Portal System, and a List (perhaps empty) of the Port IDs of
9 the operational Aggregation Ports in this Portal System.
10 DRF_Neighbor_Admin_Conversation_GatewayList_Digest
11 The value for the digest of the prioritized Gateway Conversation ID-to-Gateway assignments
12 of the Neighbor Portal System, assigned by administrator or System policy for use when the
13 Neighbor Portal System's information is unknown. Its default value is set to NULL. It is always
14 set equal to aDrniNeighborAdminConvGatewayListDigest (7.4.1.1.11).
15 Value: MD5 digest
16 DRF_Neighbor_Admin_Conversation_PortList_Digest
17 The value for the digest of the prioritized Port Conversation ID-to-Aggregation Port
18 assignments of the Neighbor Portal System, assigned by administrator or System policy for use
19 when the Neighbor Portal System's information is unknown. Its default value is set to NULL. It
20 is always set equal to aDrniNeighborAdminConvPortListDigest (7.4.1.1.12).
21 Value: MD5 digest
22 DRF_Neighbor_Admin_DRCP_State
23 Default value for the Neighbor Portal System's DRCP state parameters, assigned by
24 administrator or System policy for use when the Neighbor Portal System's information is
25 unknown or expired. The value consists of the following set of variables, as described in
26 9.4.3.2:
27 Home_Gateway
28 Neighbor_Gateway
29 Other_Gateway
30 IPP_Activity
31 DRCP_Timeout
32 Gateway_Sync
33 Port_Sync
34 Expired
35 Value: 8 bits
36 DRF_Neighbor_Admin_Gateway_Algorithm
37 The value for the gateway algorithm of the Neighbor Portal Systems, assigned by administrator
38 or System policy for use when the Neighbor Portal System's information is unknown. Its
39 default value is set to NULL. It is always set equal to aDrniNeighborAdminGatewayAlgorithm
40 (7.4.1.1.14).
41 Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining
42 this algorithm followed by one octet identifying this specific algorithm).
43 DRF_Neighbor_Admin_Port_Algorithm
44 The value for the port algorithm of the Neighbor Portal Systems, assigned by administrator or
45 System policy for use when the Neighbor Portal System's information is unknown. Its default
46 value is set to NULL. It is always set equal to aDrniNeighborAdminPortAlgorithm
47 (7.4.1.1.15).
48 Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining
49 this algorithm followed by one octet identifying this specific algorithm).
50 DRF_Portal_System_Number
51 A unique identifier for this Portal System in the Portal. Always set equal to
52 aDrniPortalSystemNumber (7.4.1.1.7). Transmitted in DRCPDUs.
53 Value: An integer in the range [1...3].
54

PSI

This variable is set to TRUE by the updateDRFHomeState function when the Portal System is isolated from the other Portal Systems within the same Portal.

Value: Boolean

9.4.9 Per-IPP variables

CC_Time_Shared

A Boolean indicating that Neighbor and Home Portal Systems on this IPP are consistently configured to use Network / IPL sharing by time.

Value: Boolean

CC_EncTag_Shared

A Boolean indicating that Neighbor and Home Portal Systems on this IPP are consistently configured to use Network / IPL sharing by tag or Network / IPL sharing by encapsulation as dictated by the Network / IPL method selected the aDrniEncapsulationMethod (7.4.1.1.17).

Value: Boolean

Differ_Conf_Portal

A Boolean indicating that the configured 14 least significant bit of administrative value of the key for the Portal Aggregator used by the immediate Neighbor Portal System on this IPP are different from the expected ones.

Value: Boolean

Differ_Conf_Portal_System_Number

A Boolean indicating the configured Portal System Numbers used by the immediate Neighbor Portal System on this IPP are different from the expected ones.

Value: Boolean

Differ_Gateway_Digest

A Boolean indicating that the Gateway_Digest used by the immediate Neighbor Portal System on this IPP is different from the expected one.

Value: Boolean

Differ_Port_Digest

A Boolean indicating that the Port_Digest used by the immediate Neighbor Portal System on this IPP is different from the expected one.

Value: Boolean

Differ_Portal

A Boolean indicating that the received DRCPDU on this IPP is associated with a different Portal.

Value: Boolean

DRF_Home_Conf_Neighbor_Portal_System_Number

This Portal System's configuration value for the Portal System Number of the Neighbor Portal System attached to this IPP. Always set equal to the value assigned to the two least significant bits of the priority component of the Port ID of this IPP (7.4.1.1.8). Transmitted in DRCPDUs.

Value: An integer in the range [1...3].

DRF_Home_Network/IPL_IPLEncap_Digest

A digest of aDrniIPLEncapMap (7.4.1.1.18), configured on this IPP, for exchange with the Neighbor Portal System on the IPL. The digest is calculated in a way similar to that specified for DRF_Home_Conversation_GatewayList_Digest where in place of the list of Portal System Numbers provided by aDrniConvAdminGateway[], the identifiers for the encapsulation method provided by aDrniIPLEncapMap are used instead. Transmitted in the Network/IPL Sharing Encapsulation TLV (9.4.3.4.2).

Value: MD5 digest

DRF_Home_Network/IPL_NetEncap_Digest

A digest of aDrniNetEncapMap (7.4.1.1.19), configured on this IPP, for exchange on the shared network link. The digest is calculated in a way similar to that specified for DRF_Home_Conversation_GatewayList_Digest where in place of the list of Portal System

1 Numbers provided by aDrniConvAdminGateway[], the identifiers for the translated values of
2 the identifiers used for a network frame provided by aDrniNetEncapMap are used instead.
3 Transmitted in the Network/IPL Sharing Encapsulation TLV (9.4.3.4.2).
4 Value: MD5 digest
5 DRF_Home_Network/IPL_Sharing_Method
6 The Network/IPL sharing method used by this DR Function to share this IPP with network
7 data. Always set equal to the aDrniEncapsulationMethod (7.4.1.1.13). Transmitted in the
8 Network/IPL Sharing Method TLV (9.4.3.4.1) when the aDrniEncapsulationMethod is not set
9 to the default NULL value.
10 Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining
11 this method followed by one octet identifying this specific method).
12 DRF_Home_Oper_DRCP_State
13 The operational values of this Portal System's DRCP state parameters as reported on this IPP.
14 This consists of the following set of variables, as described in 9.4.3.2:
15 Home_Gateway
16 Neighbor_Gateway
17 Other_Gateway
18 IPP_Activity
19 DRCP_Timeout
20 Gateway_Sync
21 Port_Sync
22 Expired
23 Value: 8 bits
24 DRF_Neighbor_Admin_Aggregator_Key
25 The administrative Aggregator Key value of the Neighbor Portal System on this IPP.
26 Transmitted in DRCPDUs.
27 Value: Integer
28 DRF_Neighbor_Aggregator_ID
29 The last received, MAC address component of Aggregator System ID of the Neighbor Portal
30 System, on this IPP.
31 Value: 48 bits
32 DRF_Neighbor_Aggregator_Priority
33 The last received, System Priority of the Neighbor Portal System's Aggregator, on this IPP.
34 Value: Integer
35 DRF_Neighbor_Conversation_GatewayList_Digest
36 The last-received Gateway Conversation ID digest of the Neighbor Portal System on this IPP.
37 Value: MD5 digest
38 DRF_Neighbor_Conversation_PortList_Digest
39 The last-received Port Conversation ID digest of the Neighbor Portal System on this IPP.
40 Value: MD5 digest
41 DRF_Neighbor_Gateway_Algorithm
42 The value of the algorithm used by the Neighbor Portal System to assign frames to Gateway
43 Conversation IDs received on this IPP.
44 Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining
45 this algorithm followed by one octet identifying this specific algorithm).
46 DRF_Neighbor_Gateway_Conversation_Mask
47 The operational value of this Portal System's view of the immediate Neighbor Portal System's
48 Gateway Vector on this IPP. This is a Boolean vector, indexed by Gateway Conversation ID,
49 indicating whether the indexed Gateway Conversation ID is enabled to pass through the
50 Neighbor Gateway (FALSE = blocked). Its value is initialized to a NULL vector and is updated
51 by the recordNeighborState function to the Gateway Vector on the first entry of the Gateway
52 Vector database for the Neighbor Portal System on this IPP.
53 Value: sequence of Boolean values, indexed by Gateway Conversation ID.
54 DRF_Neighbor_Gateway_Sequence

1 The sequence number of the operational DRF_Neighbor_Gateway_Conversation_Mask vector.
2 It is initialized to zero and is updated by the recordNeighborState function. Transmitted in
3 DRCPDUs.
4 Value: Integer

5 DRF_Neighbor_Network/IPL_IPLEncap_Digest
6 The last-received digest of aDrniIPLEncapMap of the Neighbor Portal System on this IPP.
7 Value: MD5 digest

8 DRF_Neighbor_Network/IPL_NetEncap_Digest
9 The last-received digest of aDrniNetEncapMap, for exchange on the shared network link of the
10 Neighbor Portal System on this IPP.
11 Value: MD5 digest

12 DRF_Neighbor_Network/IPL_Sharing_Method
13 The last-received Network/IPL sharing method used of the Neighbor Portal System on this IPP.
14 Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining
15 this method followed by one octet identifying this specific method).

16 DRF_Neighbor_Oper_Aggregator_Key
17 The last-received operational Aggregator Key value of the Neighbor Portal System on this IPP.
18 Value: Integer

19 DRF_Neighbor_Oper_Partner_Aggregator_Key
20 The operational Partner Aggregator Key value of the Neighbor Portal System on this IPP.
21 Transmitted in DRCPDUs.
22 Value: Integer

23 DRF_Neighbor_Oper_DRCP_State
24 The operational value of this Portal System's view of the current values of the Neighbor Portal
25 System's DRCP state parameters. The Home DR Function sets this variable to the value
26 received from the Neighbor Portal System in an DRCPDU. The value consists of the following
27 set of variables, as described in 9.4.3.2:
28 Home_Gateway
29 Neighbor_Gateway
30 Other_Gateway
31 IPP_Activity
32 DRCP_Timeout
33 Gateway_Sync
34 Port_Sync
35 Expired
36 Value: 8 bits

37 DRF_Neighbor_Conf_Portal_System_Number
38 The Neighbor Portal System's configuration Portal System Number value for this Portal
39 System that was last received on this IPP.
40 Value: An integer in the range [1...3].

41 DRF_Neighbor_Port_Algorithm
42 The value of the algorithm used by the Neighbor Portal System to assign frames to Port
43 Conversation IDs received on this IPP.
44 Value: 4-octet (3-octet OUI or CID identifying the organization that is responsible for defining
45 this algorithm followed by one octet identifying this specific algorithm).

46 DRF_Neighbor_Portal_System_Number
47 The last received identifier of the Neighbor Portal System on this IPP.
48 Value: An integer in the range [1...3].

49 DRF_Neighbor_State
50 The operational state of the immediate Neighbor Portal System on this IPP.
51 Value: Boolean flag indicating the operational state of the Neighbor Portal System's Gateway
52 (TRUE indicates operational), the Boolean Gateway Vector associated with the most recent
53 entry in the Gateway Vector database for the Neighbor Portal System, and a List (perhaps
54

1 empty) of the Port IDs of the operational Aggregation Ports of the Neighbor Portal System on
2 this IPP.

3 DRF_Other_Neighbor_Admin_Aggregator_Key
4 The administrative Aggregator Key value of the Other neighbor Portal System associated this
5 IPP. Transmitted in DRCPDUs.
6 Value: Integer

7 DRF_Other_Neighbor_Gateway_Conversation_Mask
8 The operational value of this Portal System's view of the Other neighbor Portal System's
9 Gateway Vector on this IPP. This is a Boolean vector, indexed by Gateway Conversation ID,
10 indicating whether the indexed Gateway Conversation ID is enabled to pass through the Other
11 neighbor Gateway (FALSE = blocked). Its value is initialized to a NULL vector and is updated
12 by the recordNeighborState function to the Gateway Vector on the first entry of the Gateway
13 Vector database for the Other neighbor Portal System on this IPP. Transmitted in DRCPDUs.
14 Value: sequence of Boolean values, indexed by Gateway Conversation ID.

15 DRF_Other_Neighbor_Gateway_Sequence
16 The sequence number of the operational DRF_Other_Neighbor_Gateway_Conversation_Mask
17 vector. It is initialized to zero and is updated by the recordNeighborState function. Transmitted
18 in DRCPDUs.
19 Value: Integer

20 DRF_Other_Neighbor_Oper_Partner_Aggregator_Key
21 The operational Partner Aggregator Key value of the Other neighbor Portal System associated
22 this IPP. Transmitted in DRCPDUs.
23 Value: Integer

24 DRF_Other_Neighbor_State
25 The operational state of the Other neighbor Portal System on this IPP.
26 Value: Boolean flag indicating the operational state of the Other neighbor Portal System's
27 Gateway (TRUE indicates operational), the Boolean Gateway Vector associated with the most
28 recent entry in the Gateway Vector database for the Other neighbor Portal System, and a List
29 (perhaps empty) of the Port IDs of the operational Aggregation Ports of the Other neighbor
30 Portal System on this IPP.

31 DRF_Rcv_Home_Gateway_Conversation_Mask
32 The operational value of this Portal System's Gateway Vector as reported by the Neighbor
33 Portal System on this IPP. Its value is updated by the recordNeighborState function.
34 Value: sequence of Boolean values, indexed by Gateway Conversation ID.

35 DRF_Rcv_Home_Gateway_Sequence
36 The operational value of this Portal System's Gateway sequence number as reported by the
37 Neighbor Portal System on this IPP. Its value is updated by the recordNeighborState function.
38 Value: Integer

39 DRF_Rcv_Neighbor_Gateway_Conversation_Mask
40 The operational value of the Neighbor Portal System's Gateway Vector as reported by the
41 Neighbor Portal System on this IPP itself. Its value is updated by the recordNeighborState
42 function.
43 Value: sequence of Boolean values, indexed by Gateway Conversation ID.

44 DRF_Rcv_Neighbor_Gateway_Sequence
45 The operational value of the Neighbor Portal System's Gateway sequence number as reported
46 by the Neighbor Portal System on this IPP itself. Its value is updated by the
47 recordNeighborState function.
48 Value: Integer

49 DRF_Rcv_Other_Gateway_Conversation_Mask
50 The operational value of the Other Portal System's Gateway Vector as reported by the
51 Neighbor Portal System on this IPP. Its value is updated by the recordNeighborState function.
52 Value: sequence of Boolean values, indexed by Gateway Conversation ID.

53 DRF_Rcv_Other_Gateway_Sequence
54

1 The operational value of the Other Portal System's Gateway sequence number as reported by
2 the Neighbor Portal System on this IPP. Its value is updated by the recordNeighborState
3 function.
4 Value: Integer

5 Drni_Neighbor_Common_Methods
6 The last received Common_Methods flag of the Neighbor Portal System on this IPP carried
7 within the Topology State field.
8 Value: Boolean

9 Drni_Neighbor_Gateway_Conversation
10 The last received operational Gateway_Conversation vector of the Neighbor Portal System on
11 this IPP carried within the received 2P Gateway Conversation Vector TLV (9.4.3.3.1) or the
12 received pair of 3P Gateway Conversation Vector-1 TLV (9.4.3.3.2) and 3P Gateway
13 Conversation Vector-2 TLV (9.4.3.3.3).
14 Value: a 1024-octet vector of a sequence of Portal System Numbers (0 for none), indexed by
15 Gateway Conversation ID resulting from the concatenation of the
16 3P_Gateway_Conversation_Vector-1 and the 3P_Gateway_Conversation_Vector-2 fields from
17 the received pair of 3P Gateway Conversation Vector-1 TLV and 3P Gateway Conversation
18 Vector-2 TLV, if aDrniThreePortalSystem == 1, or a 512-octet vector of a sequence of Boolean
19 values, indexed by Gateway Conversation ID copied from the 2P Gateway Conversation
20 Vector in the received 2P Gateway Conversation Vector TLV, if aDrniThreePortalSystem == 0.

21 Drni_Neighbor_Port_Conversation
22 The last received operational Port_Conversation vector of the Neighbor Portal System on this
23 IPP carried within the received 2P Port Conversation Vector TLV (9.4.3.3.4) or the received
24 pair of 3P Port Conversation Vector-1 TLV (9.4.3.3.4) and 3P Port Conversation Vector-2 TLV
25 (9.4.3.3.6).
26 Value: a 1024-octet vector of a sequence of Portal System Numbers (0 for none), indexed by
27 Port Conversation ID resulting from the concatenation of the 3P_Port_Conversation_Vector-1
28 and the 3P_Port_Conversation_Vector-2 fields from the received pair of 3P Port Conversation
29 Vector-1 TLV and 3P Port Conversation Vector-2 TLV, if aDrniThreePortalSystem == 1, or a
30 512-octet vector of a sequence of Boolean values, indexed by Port Conversation ID copied
31 from the 2P Port Conversation Vector in the received 2P Port Conversation Vector TLV, if
32 aDrniThreePortalSystem == 0.

33 Drni_Neighbor_ONN
34 The last received ONN flag of the Neighbor Portal System on this IPP carried within the
35 Topology State field.
36 Value: Boolean

37 Drni_Neighbor_Portal_Addr
38 The last received MAC address component of Portal's System ID of the Neighbor Portal
39 System on this IPP.
40 Value: 48 bits

41 Drni_Neighbor_Portal_Priority
42 The last received System Priority of the Neighbor Portal System on this IPP.
43 Value: Integer

44 Drni_Neighbor_State[]
45 The last received operational value of Drni_Portal_System_State[] used by the Neighbor Portal
46 System on this IPP. This variable is updated by the recordNeighborState function.
47 Value: For each Portal System, the Boolean flag indicating the operational state of the current
48 Portal System's Gateway (TRUE indicates operational), the Boolean Gateway vector of the
49 Portal System, and a List (perhaps empty) of the Port IDs of the operational Aggregation Ports
50 in this Portal System as reported by the Neighbor Portal System on this IPP.

51 Drni_Neighbor_Three_System_Portal
52 The last received Boolean flag indicating if the Neighbor Portal System on this IPP is a Portal
53 System that is part of a Portal consisting of three Portal Systems.
54 Value: Boolean

1 Enabled_Time_Shared
2 A Boolean indicating that Neighbor and Home Portal System on this IPP are consistently
3 configured and the Network / IPL sharing by time methods specified in 9.3.2.1 are enabled.
4 Value: Boolean

5 Enabled_EncTag_Shared
6 A Boolean indicating that Neighbor and Home Portal System on this IPP are consistently
7 configured to use the tag manipulation methods of Network / IPL sharing by tag or Network /
8 IPL sharing by encapsulation, as dictated by the Network / IPL method, selected by the
9 aDrniEncapsulationMethod (7.4.1.1.17).
10 Value: Boolean

11 Ipp_Other_Gateway_Conversation
12 The operational vector listing which Portal System (if any) is passing each Gateway
13 Conversation ID as assigned to by the immediate Neighbor Portal System on this IPP.
14 Value: sequence of Portal System Numbers (0 for none), indexed by Gateway Conversation ID.
15 Value computed from aDrniConvAdminGateway[] and Drni_Neighbor_State[] upon
16 initialization and whenever the managed object changes or IppGatewayUpdate is TRUE.

17 Ipp_Other_Port_Conversation_Portal_System
18 The operational vector listing which Portal System (if any) is passing each Port
19 Conversation ID as assigned to by the immediate Neighbor Portal System on this IPP.
20 Value: sequence of Portal System Numbers (0 for none), indexed by Port Conversation ID.
21 Value computed from Conversation_PortList[] and Drni_Neighbor_State[] upon
22 initialization and whenever the managed object changes or IppPortUpdate is TRUE.

23 IPP_port_enabled
24 A variable indicating that the IPL has been established and the IPP is operable.
25 Value: Boolean
26 TRUE if the IPP is operable (MAC_Operational == TRUE).
27 FALSE otherwise.

28
29 NOTE—The means by which the value of the IPP_port_enabled variable is generated by the underlying MAC
30 is implementation-dependent.

31
32 Ipp_Portal_System_State[]
33 The List of the states of the Portal Systems reachable through this IPP. This variable is set by
34 the updatePortalState function.
35 Value: For each Portal System, Boolean flag indicating the operational state of the current
36 Portal System's Gateway reachable through this IPP (TRUE indicates operational), the
37 associated Boolean Gateway Vector of the Portal System indicating the Gateway
38 Conversation IDs that is enabled by the network control protocol to pass through the Gateway
39 (FALSE = blocked), and a List (perhaps empty) of the Port IDs of the operational Aggregation
40 Ports in that Portal System.
41 In this list, the state of the immediately adjacent Portal System is the first state in the list.
42 The list can have at most two Portal Systems' states.

43 Missing_Rcv_Gateway_Con_Vector
44 This variable indicates that Differ_Gateway_Digest is set to TRUE and no
45 Drni_Neighbor_Gateway_Conversation can be extracted from the last received DRCPDU.
46 Value: Boolean

47 Missing_Rcv_Port_Con_Vector
48 This variable indicates that Differ_Port_Digest is set to TRUE and no
49 Drni_Neighbor_Port_Conversation can be extracted from the last received DRCPDU.
50 Value: Boolean

51 NTTDRCPDU
52 Need To Transmit flag.
53 Value: Boolean
54

1 TRUE indicates that there is new protocol information that should be transmitted on this
2 IPP, or that the Neighbor Portal System needs to be reminded of the old information.
3 FALSE otherwise.

4 ONN

5 Other Non Neighbor (ONN) flag. This value is updated by the updatePortalState function and
6 is applicable only on Portals consisting of three Portal Systems. Transmitted in DRCPDUs.

7 Value: Boolean

8 TRUE indicates that the Other Ports Information TLV is not associated with an immediate
9 Neighbor of this Portal System. FALSE (encoded as 0) indicates that the Other Ports
10 Information TLV is from an immediate Neighbor Portal System on the other IPP on this
11 Portal System.
12

13 9.4.10 Variables used for managing the operation of the state machines

14 BEGIN

15 This variable indicates the initialization (or reinitialization) of the DRCP protocol entity. It is
16 set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-
17)initialization has completed.

18 Value: Boolean

19 DRCP_Enabled

20 This variable indicates that the associated IPP is operating the DRCP. If the link is not a point-
21 to-point link, the value of DRCP_Enabled shall be FALSE. Otherwise, the value of
22 DRCP_Enabled shall be TRUE.

23 Value: Boolean

24 HomeGatewayVectorTransmit

25 This variable indicates to the DRCPDU Transmit machine (9.4.19) that it needs to transmit the
26 DRF_Home_Gateway_Conversation_Mask in the Home_Gateway_Vector field of the Home
27 Gateway Vector TLV (9.4.3.2). There is one HomeGatewayVectorTransmit variable per IPP on
28 this Portal System.

29 Value: Boolean

30 GatewayConversationTransmit

31 This variable indicates to the DRCPDU Transmit machine (9.4.19) that it needs to transmit ~~the~~
32 Gateway Conversation Vector TLVs (9.4.3.3.1, 9.4.3.3.2, 9.4.3.3.3). It is set equal to TRUE
33 when new values for the Drni_Conversation_GatewayList[] are initiated. It is set to TRUE or
34 FALSE as specified by the recordPortalConfValues function. There is one
35 GatewayConversationTransmit variable per IPP on this Portal System.

36 Value: Boolean

37 GatewayConversationUpdate

38 This variable indicates that the per Gateway Conversation ID distributions need to be updated.

39 Value: Boolean

40 IppAllGatewayUpdate

41 This variable is the logical OR of the IppGatewayUpdate variables for all IPPs in this Portal
42 System.

43 Value: Boolean

44 IppAllPortUpdate

45 This variable is the logical OR of the IppPortUpdate variables for all IPPs in this Portal System.

46 Value: Boolean

47 IppAllUpdate

48 This variable is the logical OR of the IppPortUpdate and the IppGatewayUpdate variables for
49 all IPPs in this Portal System.

50 Value: Boolean

51 IppGatewayUpdate

1 This variable indicates that the per Gateway Conversation ID distributions on the associated
2 IPP need to be updated. There is one IppGatewayUpdate variable per IPP on this Portal
3 System.
4 Value: Boolean

5 IppPortUpdate
6 This variable indicates that the per Port Conversation ID distributions on the associated IPP
7 need to be updated. There is one IppPortUpdate variable per IPP on this Portal System.
8 Value: Boolean

9 OtherGatewayVectorTransmit
10 This variable indicates to the DRCPDU Transmit machine (9.4.19) that it needs to transmit the
11 DRF_Home_Gateway_Conversation_Mask in the Other_Gateway_Vector field of the Other
12 Gateway Vector TLV (9.4.3.2). There is one OtherGatewayVectorTransmit variable per IPP on
13 this Portal System.
14 Value: Boolean

15 PortConversationTransmit
16 This variable indicates to the DRCPDU Transmit machine (9.4.19) that it needs to transmit ~~the~~
17 Port Conversation Vector TLV_s (9.4.3.3.4, [9.4.3.3.5](#), [9.4.3.3.6](#)). It is set equal to TRUE when
18 new values for the aAggConversationAdminLink[] (7.3.1.1.35) are initiate. It is set to TRUE or
19 FALSE as specified by the recordPortalConfValues function. There is one
20 PortConversationTransmit variable per IPP on this Portal System.
21 Value: Boolean

22 PortConversationUpdate
23 This variable indicates that the per Port Conversation ID distributions need to be updated.
24 Value: Boolean

25 9.4.11 Functions

26

27

28 InitializeDRNIGatewayConversation
29 This function sets the Drni_Portal_System_Gateway_Conversation to a sequence of zeros,
30 indexed by Gateway Conversation ID.

31 InitializeDRNIPortConversation
32 This function sets the Drni_Portal_System_Port_Conversation to a sequence of zeros, indexed
33 by Port Conversation ID.

34 InitializeIPPGatewayConversation
35 This function sets the Ipp_Gateway_Conversation_Direction to a sequence of zeros, indexed
36 by Gateway Conversation ID.

37 InitializeIPPPortConversation
38 This function sets the Ipp_Port_Conversation_Passes to a sequence of zeros, indexed by Port
39 Conversation ID.

40 recordDefaultDRCPDU
41 This function sets the current Neighbor Portal System's operational parameter values to the
42 default parameter values provided by the administrator as follows:
43 DRF_Neighbor_Port_Algorithm = DRF_Neighbor_Admin_Port_Algorithm;
44 DRF_Neighbor_Gateway_Algorithm = DRF_Neighbor_Admin_Gateway_Algorithm;
45 DRF_Neighbor_Conversation_PortList_Digest
46 = DRF_Neighbor_Admin_Conversation_PortList_Digest;
47 DRF_Neighbor_Conversation_GatewayList_Digest
48 = DRF_Neighbor_Admin_Conversation_GatewayList_Digest;
49 DRF_Neighbor_Oper_DRCP_State = DRF_Neighbor_Admin_DRCP_State;
50 DRF_Neighbor_Aggregator_Priority = aAggPortPartnerAdminSystemPriority (7.3.2.1.6);
51 DRF_Neighbor_Aggregator_ID = aAggPortPartnerAdminSystemID (7.3.2.1.8);
52 Drni_Neighbor_Portal_Priority = aAggPortPartnerAdminSystemPriority (7.3.2.1.6);
53 Drni_Neighbor_Portal_Addr = aAggPortPartnerAdminSystemID (7.3.2.1.8);
54 DRF_Neighbor_Portal_System_Number

1 = DRF_Home_Conf_Neighbor_Portal_System_Number, and;
2 DRF_Neighbor_Conf_Portal_System_Number = DRF_Portal_System_Number.

3
4 In addition for the Neighbor Portal System on the IPP:

5 The DRF_Neighbor_State is set to NULL (the Boolean flag for the Neighbor Portal System's
6 Gateway is set to FALSE, the Neighbor Gateway Vector is set to NULL, and the list of the
7 operational Aggregation Ports on the Neighbor Portal System on this IPP is emptied) and if
8 present, the DRF_Other_Neighbor_State is also set to NULL (the Boolean flag for the Other
9 neighbor Portal System's Gateway is set to FALSE, the Other Gateway Vector is set to NULL,
10 and the list of the operational Aggregation Ports on the Other neighbor Portal System on this
11 IPP is emptied). No Portal System state information is available for any Portal System on this
12 IPP;

13 The DRF_Neighbor_Admin_Aggregator_Key on this IPP is set to zero;

14 The DRF_Other_Neighbor_Admin_Aggregator_Key on this IPP is set to zero;

15 The DRF_Neighbor_Oper_Partner_Aggregator_Key on this IPP is set to zero;

16 The DRF_Other_Neighbor_Oper_Partner_Aggregator_Key on this IPP is set to zero;

17 The Drni_Neighbor_Gateway_Conversation on this IPP is set;

18 to All_Neighbor_Conversation, which is a Portal System Number vector with all its 4096
19 elements set to the DRF_Home_Conf_Neighbor_Portal_System_Number, if
20 Drni_Three_System_Portal == 1 or;

21 to **1**, which is a Boolean Vector with all its 4096 elements set to 1, if
22 Drni_Three_System_Portal == 0.

23 The Drni_Neighbor_Port_Conversation on this IPP is set

24 to All_Neighbor_Conversation, if Drni_Three_System_Portal == 1 or;

25 to **1**, if Drni_Three_System_Portal == 0 and;

26 The variable ChangePortal is set to TRUE.

27
28 Finally it sets CC_Time_Shared and CC_EncTag_Shared to FALSE.

29 recordNeighborState

30 This function sets DRF_Neighbor_Oper_DRCP_State.IPP_Activity to TRUE and records the
31 parameter values for the *Drni_Portal_System_State[]* and *DRF_Home_Oper_DRCP_State*
32 carried in a received DRCPDU [item s) is 9.4.3.2] on the IPP, as the current parameter values
33 for *Drni_Neighbor_State[]* and *DRF_Neighbor_Oper_DRCP_State* associated with this IPP
34 respectively. In particular, the operational Boolean Gateway Vectors for each Portal System in
35 the *Drni_Neighbor_State[]* are extracted from the received DRCPDU as follows:

36
37 For the DRF_Rcv_Neighbor_Gateway_Conversation_Mask, if the *Home_Gateway* bit in the
38 *DRF_Home_Oper_DRCP_State* carried in the received DRCPDU is 0;

39 DRF_Rcv_Neighbor_Gateway_Conversation_Mask is set to NULL;

40 Otherwise if the *Home_Gateway_Vector* field [item al) in 9.4.3.2] is present in the received
41 *Home Gateway Vector TLV* [item ai) in 9.4.3.2];

42 DRF_Rcv_Neighbor_Gateway_Conversation_Mask = *Home_Gateway_Vector*;

43 The tuple (*Home_Gateway_Sequence*, *Home_Gateway_Vector*) in the received *Home*
44 *Gateway Vector TLV* is stored as an entry in the Gateway Vector database for the Neighbor
45 Portal System on this IPP, indexed by the received *Home_Gateway_Sequence* in
46 increasing sequence number order, and;

47 The OtherGatewayVectorTransmit on the other IPP, if it exists and is operational, is set to
48 TRUE;

49 Otherwise if the *Home_Gateway_Vector* field is not present in the received *Home Gateway*
50 *Vector TLV*;

51 The OtherGatewayVectorTransmit on the other IPP, if it exists and is operational, is set to
52 FALSE, and;

53 The *Home_Gateway_Sequence* [item ak) in 9.4.3.2] is used as an index for a query in the
54 Gateway Vector database for the Neighbor Portal System on this IPP and;

1 If the tuple (*Home_Gateway_Sequence*, *Neighbor_Gateway_Vector*) is stored as the first
2 entry in the database, then;

3 DRF_Rcv_Neighbor_Gateway_Conversation_Mask = *Neighbor_Gateway_Vector*;

4 Otherwise

5 DRF_Rcv_Neighbor_Gateway_Conversation_Mask = **1**, where **1** is a Boolean
6 Vector with all its 4096 elements set to 1.

7
8 NOTE—If a Gateway bit is set but no valid Gateway Vector is available then to avoid looping the
9 recipient must assume that the neighbor Portal System has unblocked its associated gateway to all
10 conversations

11
12 For the DRF_Rcv_Home_Gateway_Conversation_Mask, if the *Neighbor_Gateway* bit in the
13 *DRF_Home_Oper_DRCP_State* carried in the received DRCPDU is 0;

14 DRF_Rcv_Home_Gateway_Conversation_Mask is set to NULL;

15 Otherwise;

16 The *Neighbor_Gateway_Sequence* [item ao) in 9.4.3.2] carried in the received *Neighbor*
17 *Gateway_Vector_TLV* [item am) in 9.4.3.2] is used as an index for a query in the Gateway
18 Vector database of this Portal System and;

19 If the tuple (*Neighbor_Gateway_Sequence*, *Home_Gateway_Vector*) is stored in the
20 database, then;

21 DRF_Rcv_Home_Gateway_Conversation_Mask = *Home_Gateway_Vector*;

22 In addition, if that is the first entry in the database, then;

23 The HomeGatewayVectorTransmit on this IPP is set to FALSE;

24 Otherwise;

25 The HomeGatewayVectorTransmit on this IPP is set to TRUE, and if the
26 *Neighbor_Gateway_Sequence* value is larger than the currently used
27 *Home_Gateway_Sequence* a new entry is created in Gateway Vector database of
28 this Portal System with the tuple values (*Neighbor_Gateway_Sequence* + 1,
29 *Home_Gateway_Vector*);

30 Otherwise

31 DRF_Rcv_Home_Gateway_Conversation_Mask = **1**, where **1** is a Boolean Vector
32 with all its 4096 elements set to 1, and;

33 The HomeGatewayVectorTransmit is set to TRUE.

34
35 For the DRF_Rcv_Other_Gateway_Conversation_Mask, if the *Other_Gateway* bit in the
36 *DRF_Home_Oper_DRCP_State* carried in the received DRCPDU is 0;

37 DRF_Rcv_Other_Gateway_Conversation_Mask is set to NULL;

38 Otherwise if the *Other_Gateway_Vector* field [item as) in 9.4.3.2] is present in the received
39 *Other_Gateway_Vector_TLV* [item ap) in 9.4.3.2];

40 DRF_Rcv_Other_Gateway_Conversation_Mask = *Other_Gateway_Vector*; and

41 If on this IPP, *Drni_Neighbor_ONN* == FALSE;

42 The tuple (*Other_Gateway_Sequence*, *Other_Gateway_Vector*) in the received *Other*
43 *Gateway_Vector_TLV* is stored as an entry in the Gateway Vector database for the
44 Other neighbor Portal System on this IPP indexed by the received
45 *Other_Gateway_Sequence* in increasing sequence number order;

46 Otherwise if the *Other_Gateway_Vector* field is not present in the received *Other_Gateway*
47 *Vector_TLV*;

48 The *Other_Gateway_Sequence* [item ar) in 9.4.3.2] is used as an index for a query in the
49 Gateway Vector database for the Other neighbor Portal System on this IPP and;

50 If the tuple (*Other_Gateway_Sequence*, *Other_Gateway_Vector*) is stored in the database,
51 then;

52 DRF_Rcv_Other_Gateway_Conversation_Mask = *Other_Gateway_Vector*;

53 In addition, if that is the first entry in the database, then;

54 The OtherGatewayVectorTransmit on this IPP is set to FALSE;

1 Otherwise;

2 The OtherGatewayVectorTransmit on this IPP is set to TRUE;

3 Otherwise

4 DRF_Rcv_Other_Gateway_Conversation_Mask = **1**, where **1** is a Boolean Vector
5 with all its 4096 elements set to 1, and;

6 The OtherGatewayVectorTransmit on this IPP is set to TRUE.

7
8 It also records the variables below as follows:

9 The parameter values for the *Home_Gateway* in the *DRF_Home_Oper_DRCP_State*, the
10 Gateway Vector from the most recent entry in the Gateway Vector database for the Neighbor
11 Portal System, and the *Active_Home_Ports* in the *Home Ports Information TLV*, carried in a
12 received DRCPDU on the IPP, are used as the current values for the DRF_Neighbor_State on
13 this IPP and are associated with the Portal System identified by
14 DRF_Neighbor_Portal_System_Number;

15 The parameter values for the *Other_Gateway* in the *DRF_Home_Oper_DRCP_State*, the
16 Gateway Vector from the most recent entry in the Gateway Vector database for the Other
17 neighbor Portal System, and the *Other_Neighbor_Ports* in the *Other Ports Information TLV*,
18 carried in a received DRCPDU on the IPP, are used as the current values for the
19 DRF_Other_Neighbor_State on this IPP and are associated with the Portal System identified by
20 the value assigned to the two most significant bits of the
21 *DRF_Other_Neighbor_Admin_Aggregator_Key* carried within the Other Ports Information
22 TLV in the received DRCPDU. If no Other Ports Information TLV is carried in the received
23 DRCPDU and the Portal Topology contains three Portal Systems, the
24 DRF_Other_Neighbor_State is set to NULL (Other_Gateway is set to FALSE, the
25 Other_Gateway_Vector is set to NULL, and the list of the operational Aggregation Ports on the
26 Other neighbor Portal System on this IPP is emptied) and no Portal System state information is
27 available on this IPP for the distant Neighbor Portal System on the IPP;

28 DRF_Neighbor_Admin_Aggregator_Key = *DRF_Home_Admin_Aggregator_Key*;

29 DRF_Neighbor_Oper_Partner_Aggregator_Key

30 = *DRF_Home_Oper_Partner_Aggregator_Key*;

31 DRF_Other_Neighbor_Admin_Aggregator_Key

32 = *DRF_Other_Neighbor_Admin_Aggregator_Key*, and;

33 DRF_Other_Neighbor_Oper_Partner_Aggregator_Key

34 = *DRF_Other_Neighbor_Oper_Partner_Aggregator_Key*.

35 Both DRF_Other_Neighbor_Admin_Aggregator_Key and
36 DRF_Other_Neighbor_Oper_Partner_Aggregator_Key are set to NULL when the received
37 DRCPDU does not contain the Other Ports Information TLV [item ad) in 9.4.3.2].

38
39 In addition, if Network / IPL sharing by time (9.3.2.1) is supported, the function records the
40 parameter value for the *DRF_Home_Network/IPL_Sharing_Method* carried in the received
41 Network/IPL Sharing Method TLV (9.4.3.4.1) as the current parameter value for the
42 DRF_Neighbor_Network/IPL_Sharing_Method and if this is the same as the System's
43 DRF_Home_Network/IPL_Sharing_Method, it sets CC_Time_Shared to TRUE, otherwise it
44 sets CC_Time_Shared to FALSE.

45
46 Further, if Network / IPL sharing by tag (9.3.2.2) or Network / IPL sharing by encapsulation
47 (9.3.2.3) is supported, the function records the Neighbor Portal System's Network/IPL sharing
48 related parameter values carried in the received Network/IPL sharing TLVs (9.4.3.4) from an
49 IPP, as the current operational parameter values for the immediate Neighbor Portal System on
50 this IPP as follows:

51 DRF_Neighbor_Network/IPL_Sharing_Method

52 = *DRF_Home_Network/IPL_Sharing_Method*, carried in the received Network/IPL Sharing
53 Method TLV (9.4.3.4.1);

54 DRF_Neighbor_Network/IPL_IPLEncap_Digest

1 = *DRF_Home_Network/IPL_IPLEncap_Digest*, carried in the received Network/IPL Sharing
2 Encapsulation TLV (9.4.3.4.2); and
3 *DRF_Neighbor_Network/IPL_NetEncap_Digest*
4 = *DRF_Home_Network/IPL_NetEncap_Digest* carried in the received Network/IPL Sharing
5 Encapsulation TLV (9.4.3.4.2).
6 It then compares the newly updated values of the Neighbor Portal System to this Portal
7 System's expectations and if
8 *DRF_Neighbor_Network/IPL_Sharing_Method*
9 == *DRF_Home_Network/IPL_Sharing_Method*, and
10 *DRF_Neighbor_Network/IPL_IPLEncap_Digest*
11 == *DRF_Home_Network/IPL_IPLEncap_Digest*, and
12 *DRF_Neighbor_Network/IPL_NetEncap_Digest*
13 == *DRF_Home_Network/IPL_NetEncap_Digest*, then
14 it sets *CC_EncTag_Shared* to TRUE;
15 Otherwise if one or more of the comparisons shows that the values differ,
16 it sets *CC_EncTag_Shared* to FALSE.
17
18 It then compares the Gateway operational state and Gateway Vector for each Portal System as
19 reported by this Portal System's *Drni_Portal_System_State[]* to the Gateway operational state
20 and Gateway Vector for the same Portal System as reported by the *Drni_Neighbor_State[]* and
21 if any of these differ;
22 It sets *GatewayConversationUpdate* to TRUE;
23 It sets *DRF_Home_Oper_DRCP_State.Gateway_Sync* to FALSE and;
24 If *Missing_Rcv_Gateway_Con_Vector* == TRUE;
25 The *Drni_Neighbor_Gateway_Conversation* on this IPP is set;
26 to *All_Neighbor_Conversation*, which is a Portal System Number vector with
27 all its 4096 elements set to the
28 *DRF_Home_Conf_Neighbor_Portal_System_Number*, if
29 *Drni_Three_System_Portal* == 1 or;
30 to **1**, which is a Boolean Vector with all its 4096 elements set to 1, if
31 *Drni_Three_System_Portal* == 0;
32 Otherwise;
33 *Drni_Neighbor_Gateway_Conversation* remains unchanged;
34 Otherwise if they are the same and *DRF_Home_Oper_DRCP_State.Gateway_Sync* is FALSE;
35 It sets *GatewayConversationUpdate* to TRUE and;
36 *DRF_Home_Oper_DRCP_State.Gateway_Sync* is set to TRUE;
37 Otherwise if they are the same and *Differ_Gateway_Digest* == TRUE;
38 It sets *GatewayConversationUpdate* to TRUE;
39 Otherwise;
40 *GatewayConversationUpdate* remains unchanged and;
41 *DRF_Home_Oper_DRCP_State.Gateway_Sync* is set to TRUE.
42
43 It also compares the list of the Port IDs of the operational Aggregation Ports for each Portal
44 System as reported by this Portal System's *Drni_Portal_System_State[]* to list of the Port IDs
45 of the operational Aggregation Ports for the same Portal Systems as reported by the
46 *Drni_Neighbor_State[]* and
47 if any of these differ;
48 It sets *PortConversationUpdate* to TRUE and;
49 *DRF_Home_Oper_DRCP_State.Port_Sync* is set to FALSE and;
50 If *Missing_Rcv_Port_Con_Vector* == TRUE;
51 The *Drni_Neighbor_Port_Conversation* on this IPP is set
52 to *All_Neighbor_Conversation*, if *Drni_Three_System_Portal* == 1 or;
53 to **1**, if *Drni_Three_System_Portal* == 0;
54 Otherwise;

1 Drni_Neighbor_Port_Conversation remains unchanged;
2 Otherwise if they are the same and Differ_Port_Digest == TRUE
3 It sets PortConversationUpdate to TRUE;
4 Otherwise if they are the same and DRF_Home_Oper_DRCP_State.Port_Sync is FALSE;
5 It sets PortConversationUpdate to TRUE and;
6 DRF_Home_Oper_DRCP_State.Port_Sync is set to TRUE;
7 Otherwise;
8 PortConversationUpdate remains unchanged and
9 DRF_Home_Oper_DRCP_State.Port_Sync is set to TRUE.

10 recordPortalConfValues

11 This function records the Neighbor Portal System's values carried in the *Portal Configuration*
12 *Information TLV* of a received DRCPDU (9.4.3.2) from an IPP, as the current operational
13 parameter values for the immediate Neighbor Portal System on this IPP as follows:

14 DRF_Neighbor_Portal_System_Number = *DRF_Portal_System_Number*;
15 DRF_Neighbor_Conf_Portal_System_Number
16 = *DRF_Home_Conf_Neighbor_Portal_System_Number*;
17 Drni_Neighbor_Three_System_Portal = *Drni_Three_System_Portal*;
18 Drni_Neighbor_Common_Methods = *Drni_Common_Methods*;
19 Drni_Neighbor_ONN = *ONN*;
20 DRF_Neighbor_Oper_Aggregator_Key = *DRF_Home_Oper_Aggregator_Key*;
21 DRF_Neighbor_Port_Algorithm = *DRF_Home_Port_Algorithm*;
22 DRF_Neighbor_Conversation_PortList_Digest = *DRF_Home_Conversation_PortList_Digest*;
23 DRF_Neighbor_Gateway_Algorithm = *DRF_Home_Gateway_Algorithm*; and
24 DRF_Neighbor_Conversation_GatewayList_Digest
25 = *DRF_Home_Conversation_GatewayList_Digest*.

26
27
28 It then compares the newly updated values of the Neighbor Portal System to this Portal
29 System's expectations and if the comparison of

30 DRF_Neighbor_Portal_System_Number
31 to DRF_Home_Conf_Neighbor_Portal_System_Number, or
32 DRF_Neighbor_Conf_Portal_System_Number to DRF_Portal_System_Number, or
33 Drni_Neighbor_Three_System_Portal to Drni_Three_System_Portal, or
34 Drni_Neighbor_Common_Methods to Drni_Common_Methods, or
35 the 14 least significant bits of DRF_Neighbor_Oper_Aggregator_Key to
36 the 14 least significant bits of DRF_Home_Oper_Aggregator_Key, or
37 DRF_Neighbor_Port_Algorithm to DRF_Home_Port_Algorithm, or
38 DRF_Neighbor_Conversation_PortList_Digest
39 to DRF_Home_Conversation_PortList_Digest, or
40 DRF_Neighbor_Gateway_Algorithm to DRF_Home_Gateway_Algorithm, or
41 DRF_Neighbor_Conversation_GatewayList_Digest
42 to DRF_Home_Conversation_GatewayList_Digest show that one or more of the compared
43 pairs differ;

44 The associated pairs of variables having the different values are stored in
45 aIPPDebugDifferPortalReason (7.4.4.1.4); and

46 The reportToManagement function is invoked to report to the Management system the
47 associated differences.

48 If only any of the first two pairs differ;

49 The variable Differ_Conf_Portal_System_Number is set to TRUE;

50 Otherwise;

51 The variable Differ_Conf_Portal_System_Number is set to FALSE.

52
53 In addition if the 14 least significant bits of DRF_Neighbor_Oper_Aggregator_Key ==
54 the 14 least significant bits of DRF_Home_Oper_Aggregator_Key;

1 the variable Differ_Conf_Portal is set to FALSE, and;
2 if DRF_Neighbor_Oper_Aggregator_Key is different than the
3 DRF_Home_Oper_Aggregator_Key;
4 ChangePortal will be set to TRUE;
5 Otherwise;
6 The variable Differ_Conf_Portal is set to TRUE.

7
8 Further, if the variable Differ_Conf_Portal is set to FALSE and one or more of the comparisons
9 Drni_Neighbor_Three_System_Portal to Drni_Three_System_Portal, or
10 DRF_Neighbor_Gateway_Algorithm to DRF_Home_Gateway_Algorithm differ;

11 The Drni_Neighbor_Gateway_Conversation on this IPP is set:

12 to All_Neighbor_Conversation, which is a Portal System Number vector with all its
13 4096 elements set to the DRF_Home_Conf_Neighbor_Portal_System_Number, if
14 Drni_Three_System_Portal == 1 or;
15 to **1**, which is a Boolean Vector with all its 4096 elements set to 1, if
16 Drni_Three_System_Portal == 0;

17 The variable Differ_Gateway_Digest is set to TRUE;

18 Otherwise if the variable Differ_Conf_Portal is set to FALSE and;

19 Drni_Neighbor_Three_System_Portal == Drni_Three_System_Portal and;

20 DRF_Neighbor_Gateway_Algorithm == DRF_Home_Gateway_Algorithm and;

21 If DRF_Neighbor_Conversation_GatewayList_Digest

22 == DRF_Home_Conversation_GatewayList_Digest;

23 The variable Differ_Gateway_Digest is set to FALSE;

24 The variable GatewayConversationTransmit is set to FALSE and;

25 The variable Missing_Rcv_Gateway_Con_Vector is set to FALSE;

26 Otherwise if the comparison shows that the digests differ;

27 The variable Differ_Gateway_Digest is set to TRUE;

28 The variable GatewayConversationTransmit is set to TRUE;

29 The pair of variables is stored in aIPPDebugDifferPortalReason (7.4.4.1.4); and

30 if the 2P Gateway Conversation Vector TLV (9.4.3.3.1) is present in the received
31 DRCPDU and aDrniThreePortalSystem is FALSE the TLV's length field
32 corresponds to Portal System's expectations (1024 octets when
33 aDrniThreePortalSystem is TRUE and 512 octets otherwise);

34 The Neighbor Portal System's
35 ~~Drni_Gateway_Conversation~~2P_Gateway_Conversation_Vector carried in the
36 2P Gateway Conversation Vector TLV, is recorded as the current operational
37 parameter value for the immediate Neighbor Portal System on this IPP as
38 follows:

39 Drni_Neighbor_Gateway_Conversation =

40 ~~Gateway_Conversation~~2P_Gateway_Conversation_Vector and;

41 The variable Missing_Rcv_Gateway_Con_Vector is set to FALSE;

42 Otherwise if the 3P Gateway Conversation Vector-1 TLV (9.4.3.3.2) and 3P Gateway
43 Conversation Vector-2 TLV (9.4.3.3.3) are present in the received DRCPDU and
44 aDrniThreePortalSystem is TRUE;

45 The variable 3P_Gateway_Conversation_Vector created from the concatenation
46 of the Neighbor Portal System's 3P_Gateway_Conversation_Vector-1 carried
47 in the 3P_Gateway_Conversation_Vector-1 TLV with the Neighbor Portal
48 System's 3P_Gateway_Conversation_Vector-2 carried in the 3P_Gateway
49 Conversation Vector-2 TLV, is recorded as the current operational parameter
50 value for the immediate Neighbor Portal System on this IPP as follows:

51 Drni_Neighbor_Gateway_Conversation = 3P_Gateway_Conversation_Vector
52 and;

53 The variable Missing_Rcv_Gateway_Con_Vector is set to FALSE;

Otherwise if ~~the no 2P Gateway Conversation Vector TLV(9.4.3.3.1) Gateway Conversation Vector TLVs (9.4.3.3.1, 9.4.3.3.2, 9.4.3.3.3) is are not~~ present in the received DRCPDU, $\text{Drni_Neighbor_Common_Methods} == \text{Drni_Common_Methods} == \text{TRUE}$ and the ~~2P Port Conversation Vector TLV(9.4.3.3.4) Port Conversation Vector TLVs (9.4.3.3.4, 9.4.3.3.4, 9.4.3.3.6)~~ associated with the expected number of Portal Systems in the Portal ~~is are~~ present in the received DRCPDU ~~and the 2P Port Conversation Vector TLV's length field corresponds to Portal System's expectations (1024 octets when aDrniThreePortalSystem is TRUE and 512 octets otherwise) then;~~

If $\text{aDrniThreePortalSystem} == \text{FALSE}$,

$\text{Drni_Neighbor_Gateway_Conversation} =$

~~2P Port Conversation Vector~~ $\text{Drni_Neighbor_Port_Conversation}$;

Otherwise if $\text{aDrniThreePortalSystem} == \text{TRUE}$,

$\text{Drni_Neighbor_Gateway_Conversation} =$

concatenation of $\text{3P_Port_Conversation_Vector-1}$ with

$\text{3P_Port_Conversation_Vector-2}$;

and;

The variable $\text{Missing_Rcv_Gateway_Con_Vector}$ is set to FALSE;

Otherwise;

The variable $\text{Drni_Neighbor_Gateway_Conversation}$ remains unchanged and;

The variable $\text{Missing_Rcv_Gateway_Con_Vector}$ is set to TRUE.

Finally, if the variable $\text{Differ_Conf_Portal}$ is set to FALSE and one or more of the comparisons $\text{Drni_Neighbor_Three_System_Portal}$ to $\text{Drni_Three_System_Portal}$, or $\text{DRF_Neighbor_Port_Algorithm}$ to $\text{DRF_Home_Port_Algorithm}$ differ;

The $\text{Drni_Neighbor_Port_Conversation}$ on this IPP is set:

to $\text{All_Neighbor_Conversation}$, which is a Portal System Number vector with all its 4096 elements set to the $\text{DRF_Home_Conf_Neighbor_Portal_System_Number}$, if $\text{Drni_Three_System_Portal} == 1$ or;

to $\mathbf{1}$, which is a Boolean Vector with all its 4096 elements set to 1, if $\text{Drni_Three_System_Portal} == 0$;

The variable $\text{Differ_Port_Digest}$ is set to TRUE;

Otherwise if the variable $\text{Differ_Conf_Portal}$ is set to FALSE and;

$\text{Drni_Neighbor_Three_System_Portal} == \text{Drni_Three_System_Portal}$ and;

$\text{DRF_Neighbor_Port_Algorithm} == \text{DRF_Home_Port_Algorithm}$ and;

If $\text{DRF_Neighbor_Conversation_PortList_Digest}$

$== \text{DRF_Home_Conversation_PortList_Digest}$;

The variable $\text{Differ_Port_Digest}$ is set to FALSE and;

The variable $\text{PortConversationTransmit}$ is set to FALSE;

The variable $\text{Missing_Rcv_Port_Con_Vector}$ is set to FALSE;

Otherwise if the comparison shows that the digests differ;

The variable $\text{Differ_Port_Digest}$ is set to TRUE;

The variable $\text{PortConversationTransmit}$ is set to TRUE;

~~The pair of variables is stored in aIPPDebugDifferPortalReason (7.4.4.1.4);~~

~~The reportToManagement function is invoked to report to the Management system the difference; and~~

if the 2P Port Conversation Vector TLV(9.4.3.3.4) is present in the received DRCPDU ~~and aDrniThreePortalSystem is FALSE;~~

The Neighbor Portal System's

~~Drni Port Conversation~~ $\text{2P_Port_Conversation_Vector}$ carried in the $\text{2P_Port_Conversation_Vector_TLV}$, is recorded as the current operational parameter value for the immediate Neighbor Portal System on this IPP as follows:

$\text{Drni_Neighbor_Port_Conversation} =$

~~Drni Port Conversation~~ $\text{2P_Port_Conversation_Vector}$ and;

1 The variable `Missing_Rcv_Port_Con_Vector` is set to FALSE;
 2 Otherwise if the 3P Port Conversation Vector-1 TLV (9.4.3.3.5) and 3P Port
 3 Conversation Vector-2 TLV (9.4.3.3.6) are present in the received DRCPDU and
 4 aDrniThreePortalSystem is TRUE;
 5 The variable 3P Port Conversation Vector created from the concatenation of
 6 the Neighbor Portal System's 3P Port Conversation Vector-1 carried in the 3P
 7 Port Conversation Vector-1 TLV with the Neighbor Portal System's
 8 3P Port Conversation Vector-2 carried in the 3P Port Conversation Vector-2
 9 TLV, is recorded as the current operational parameter value for the immediate
 10 Neighbor Portal System on this IPP as follows:
 11 Drni_Neighbor_Port_Conversation = 3P Port Conversation Vector and;
 12 The variable `Missing_Rcv_Port_Con_Vector` is set to FALSE;
 13 Otherwise if ~~the no 2P Port Conversation Vector TLV(9.4.3.3.4)~~ Port Conversation
 14 Vector TLVs (9.4.3.3.4, 9.4.3.3.5, 9.4.3.3.6) isare-not present in the received
 15 DRCPDU, `Drni_Neighbor_Common_Methods == Drni_Common_Methods ==`
 16 TRUE and the ~~2P Gateway Conversation Vector TLV(9.4.3.3.1)~~ Port Conversation
 17 Vector TLVs (9.4.3.3.4, 9.4.3.3.5, 9.4.3.3.6) is associated with the expected number of
 18 Portal Systems in the Portal are present in the received DRCPDU then;
 19 If `aDrniThreePortalSystem == FALSE`,
 20 Drni_Neighbor_Port_Conversation = 2P Gateway Conversation Vector;
 21 Otherwise if `aDrniThreePortalSystem == TRUE`,
 22 Drni_Neighbor_Port_Conversation =
 23 concatenation of 3P Gateway Conversation Vector-1 with
 24 3P Gateway Conversation Vector-2;
 25 Drni_Neighbor_Gateway_Conversation
 26 and;
 27 The variable `Missing_Rcv_Port_Con_Vector` is set to FALSE;
 28 Otherwise;
 29 The variable `Drni_Neighbor_Port_Conversation` remains unchanged and;
 30 The variable `Missing_Rcv_Port_Con_Vector` is set to TRUE.

31
 32 In all the operations above when the `Drni_Neighbor_Gateway_Conversation` or
 33 `Drni_Neighbor_Port_Conversation` are extracted from the Conversation Vector field in a
 34 received ~~1024-octet~~ Conversation Vector TLV, the `Drni_Neighbor_Gateway_Conversation` or
 35 `Drni_Neighbor_Port_Conversation` will be set to `All_Neighbor_Conversation`, if
 36 `Differ_Conf_Portal_System_Number == TRUE`.

37 recordPortalValues

38 This function records the Neighbor's Portal parameter values carried in a received DRCPDU
 39 (9.4.3.2) from an IPP, as the current operational parameter values for the immediate Neighbor
 40 Portal System on this IPP, as follows:

41 `DRF_Neighbor_Aggregator_Priority = Drni_Aggregator_Priority;`
 42 `DRF_Neighbor_Aggregator_ID = Drni_Aggregator_ID;`
 43 `Drni_Neighbor_Portal_Priority = Drni_Portal_Priority`, and;
 44 `Drni_Neighbor_Portal_Addr = Drni_Portal_Addr.`

45
 46 It then compares the newly updated values of the Neighbor Portal System to this Portal
 47 System's expectations and if

48 `DRF_Neighbor_Aggregator_Priority == Drni_Aggregator_Priority` and
 49 `DRF_Neighbor_Aggregator_ID == Drni_Aggregator_ID` and
 50 `Drni_Neighbor_Portal_Priority == Drni_Portal_Priority` and
 51 `Drni_Neighbor_Portal_Addr == Drni_Portal_Addr` then,
 52 the variable `Differ_Portal` is set to FALSE;
 53 Otherwise if one or more of the comparisons shows that the values differ,
 54

1 the variable Differ_Portal is set to TRUE and the associated set of variables having the
2 different values are available in aIPPDebugDifferPortalReason (7.4.4.1.4) and are
3 reported to the Management system by invoking the reportToManagement function.

4 reportToManagement

5 This function alerts the Management system of the potential existence of a Portal System
6 configuration error in this Portal due to the receipt of a misconfigured DRCPDU and sends to it
7 the conflicting information from the misconfigured received DRCPDU.

8 setDefaultPortalSystemParameters

9 This function sets this Portal System's variables to administrative set values as follows:

10 Drni_Aggregator_Priority = aAggActorSystemPriority (7.3.1.1.5);

11 Drni_Aggregator_ID = aAggActorSystemID (7.3.1.1.4);

12 Drni_Portal_Priority = aDrniPortalPriority (7.4.1.1.5);

13 Drni_Portal_Addr = aDrniPortalAddr (7.4.1.1.4);

14 DRF_Portal_System_Number = aDrniPortalSystemNumber (7.4.1.1.7);

15 DRF_Home_Admin_Aggregator_Key = aAggActorAdminKey (7.3.1.1.7);

16 DRF_Home_Port_Algorithm = aAggPortAlgorithm (7.3.1.1.33);

17 DRF_Home_Gateway_Algorithm = aDrniGatewayAlgorithm (7.4.1.1.13);

18 DRF_Home_Conversation_PortList_Digest = the MD5 digest on
19 aAggConversationAdminLink[] (7.3.1.1.35);

20 DRF_Home_Conversation_GatewayList_Digest = the MD5 digest on
21 aDrniConvAdminGateway[] (7.4.1.1.10), and;

22 DRF_Home_Oper_DRCP_State = DRF_Neighbor_Admin_DRCP_State.

23

24 In addition, it sets the Drni_Portal_System_State[] as if all Gateways in the Portal are reported
25 as FALSE, the Gateway Vectors are reported as NULL and no Aggregation Port on any Portal
26 System is reported as operational, and each Portal System Gateway Vector database is set to
27 contain a single entry having the Gateway Vector set to NULL and the associated Gateway
28 Sequence initialized to 0.

29 setGatewayConversation

30 This function sets Drni_Gateway_Conversation to the values computed from
31 aDrniConvAdminGateway[] and the current Drni_Portal_System_State[] as follows:

32 For every indexed Gateway Conversation ID, a Portal System Number is identified by
33 choosing the highest priority Portal System Number in the list of Portal System Numbers
34 provided by aDrniConvAdminGateway[] when only the Portal Systems having that Gateway
35 Conversation ID enabled in the Gateway Vectors of the Drni_Portal_System_State[] variable,
36 are included.

37 setIPPGatewayConversation

38 This function sets Ipp_Other_Gateway_Conversation as follows:

39 If Differ_Gateway_Digest == TRUE and Drni_Three_System_Portal == 1;

40 Ipp_Other_Gateway_Conversation = Drni_Neighbor_Gateway_Conversation;

41 Otherwise if Differ_Gateway_Digest == TRUE and Drni_Three_System_Portal == 0;

42 The Ipp_Other_Gateway_Conversation is extracted from the
43 Drni_Neighbor_Gateway_Conversation as follows:

44 For every indexed Gateway Conversation ID in the
45 Drni_Neighbor_Gateway_Conversation Boolean vector, the value zero bits are replaced
46 with the DRF_Portal_System_Number and the value one bits are replaced with the
47 DRF_Home_Conf_Neighbor_Portal_System_Number;

48 Otherwise if Differ_Gateway_Digest == FALSE;

49 This function sets Ipp_Other_Gateway_Conversation to the values computed from
50 aDrniConvAdminGateway[] and the Drni_Neighbor_State[] as follows:

51 For every indexed Gateway Conversation ID, a Portal System Number is identified by
52 choosing the highest priority Portal System Number in the list of Portal System Numbers
53 provided by aDrniConvAdminGateway[] when only the Portal Systems having that
54

1 Gateway Conversation ID enabled in the Gateway Vectors of the Drni_Neighbor_State[]
2 variable, are included.

3 setIPPGatewayUpdate
4 This function sets the IppGatewayUpdate on every IPP on this Portal System to TRUE.

5 setIPPPortConversation
6 This function sets Ipp_Other_Port_Conversation_Portal_System as follows:
7 If Differ_Port_Digest == TRUE and Drni_Three_System_Portal == 1;
8 Ipp_Other_Port_Conversation_Portal_System = Drni_Neighbor_Port_Conversation;
9 Otherwise if Differ_Port_Digest == TRUE and Drni_Three_System_Portal == 0;
10 The Ipp_Other_Port_Conversation_Portal_System is extracted from the
11 Drni_Neighbor_Port_Conversation as follows:
12 For every indexed Port Conversation ID in the Drni_Neighbor_Port_Conversation
13 Boolean vector, the value zero bits are replaced with the DRF_Portal_System_Number
14 and the value one bits are replaced with the
15 DRF_Home_Conf_Neighbor_Portal_System_Number;
16 Otherwise if Differ_Port_Digest == FALSE;
17 This function sets Ipp_Other_Port_Conversation_Portal_System to the values computed from
18 Conversation_PortList[] and the Drni_Neighbor_State[] as follows:
19 For every indexed Port Conversation ID, a Portal System Number is identified by choosing the
20 highest priority Portal System Number in the list of Portal Systems Numbers provided by
21 Conversation_PortList[] when only the operational Aggregation Ports, as provided by the
22 associated Lists of the Drni_Neighbor_State[] variable, are included.

23 setIPPPortUpdate
24 This function sets the IppPortUpdate on every IPP on this Portal System to TRUE.

25 setPortConversation
26 This function sets Drni_Port_Conversation to the values computed from
27 Conversation_PortList[] and the current Drni_Portal_System_State[] as follows:
28 For every indexed Port Conversation ID, a Portal System Number is identified by
29 extracting the least significant two bits of the priority component of the highest priority
30 Port ID (6.3.4) in the list of Port IDs provided by Conversation_PortList[] when only the
31 operational Aggregation Ports, as provided by the associated Lists of the
32 Drni_Portal_System_State[] variable, are included.

33 updateDRFHomeState
34 This function updates the DRF_Home_State based on the operational state of the local ports as
35 follows:
36 It sets the Home Gateway bit to TRUE or FALSE based on the mechanisms that are used to
37 identify the operational state of the local Gateway (TRUE indicates operable (i.e. the local DR
38 Function is able to relay traffic through its Gateway Port and at least one of its other Ports -
39 IPP(s) or Aggregator) and that connectivity through the local Gateway is enabled by the
40 operation of the network control protocol). In addition the operational Boolean Gateway
41 Vector, indexed by Gateway Conversation ID, is used to indicate which individual Gateway
42 Conversation IDs are enabled by the network control protocol to pass through the Home
43 Gateway (FALSE = blocked).

44
45 The operation above where the Boolean Gateway Vector is set by the network control protocol
46 respects the requirements on distribution independence and fault isolation so that the frame
47 distribution algorithm that is used to satisfy network requirements can be different from the
48 algorithm used to assign frames to the Aggregation Ports of a Link Aggregation Group and that
49 faults on the Aggregation Ports do not affect the attached network operation. On the other hand,
50 the Gateway algorithm and the Port algorithm can use the same means for assigning a frame to
51 a Conversation ID, so that the Gateway Conversation ID equals the Port Conversation ID to
52 allow matching a gateway to the link carrying a given conversation in order to minimize traffic
53 on the IPLs. In this case the Boolean Gateway Vector is not controlled by the network control
54 protocol but is set instead to the Drni_Portal_System_Port_Conversation. Additionally, this

1 function sets the Home Gateway bit TRUE if the Home Gateway Vector is not NULL and the
2 local Gateway is operable (i.e. the local DR Function is able to relay traffic through its Gateway
3 Port and its Aggregator Port). The profile of operation is configured by the
4 aDrniPortConversationControl (7.4.1.1.23) managed object.
5

6 The current operational Boolean Gateway Vector along with its associated Gateway Sequence
7 number is stored as a (Home Gateway Sequence, Home Gateway Vector) tuple in the Gateway
8 Vector database for this Portal System. If there is any change in the Home Gateway Vector, the
9 Home Gateway Sequence number of the current first entry in the Gateway Vector database for
10 this Portal System (Home Gateway Sequence, Home Gateway Vector) is increased by one and
11 a new first entry tuple (Home Gateway Sequence +1, New Home Gateway Vector) is created
12 for the New Home Gateway Vector on top of the current one and
13 HomeGatewayVectorTransmit is set to TRUE. When the Home_Gateway bit is set to FALSE,
14 the New Home Gateway Vector is set to NULL (a 4096 Boolean Vector that has all its elements
15 set to 0).
16

17 NOTE—Use of sequence numbers allows the receiver of a DRCPDU to test for a new vector without
18 exhaustive comparison of vector values.
19

20 The list of operational Aggregation Ports is created by including only those Aggregation Port
21 IDs for which the attached Aggregator reports them as having
22 Actor_Oper_Port_State.Distributing == TRUE (condition that excludes the cases where the
23 associated Aggregation Ports are either non operable (port_enabled = FALSE), in an EXPIRED
24 state, or not in the LAG) and;

25 The PSI is set to TRUE if DRF_Neighbor_Oper_DRCP_State.IPP_Activity == FALSE on all
26 IPPs on this Portal System, otherwise PSI is set to FALSE.

27 In addition, if PSI == TRUE and Home Gateway == FALSE then
28 Actor_Oper_Port_State.Synchronization is set to FALSE on all Aggregation Ports on this
29 Portal System.
30

31 The function also sets:

32 GatewayConversationUpdate to TRUE if the operational state of Gateway or the
33 configured lists for Drni_Conversation_GatewayList[] has changed and sets
34 PortConversationUpdate to TRUE if there has been any change in the list of the
35 operational Aggregation Ports as reported by changes in the associated
36 Actor_Oper_Port_State.Distributing variables or the configured lists for the
37 aAggConversationAdminLink[], otherwise;

38 GatewayConversationUpdate and PortConversationUpdate remain unchanged.
39

updateIPPGatewayConversationDirection

40 This function computes a value for Ipp_Gateway_Conversation_Direction as follows:

41 For each Gateway Conversation ID, the value is TRUE if and only if:

- 42 a) the variables Drni_Gateway_Conversation and Ipp_Portal_System_State[] indicate that
- 43 the target Portal System for this Gateway Conversation ID lies behind this IPP, and
- 44 b) Drni_Gateway_Conversation and Ipp_Other_Gateway_Conversation are in agreement
- 45 as to which Portal System should get this Gateway Conversation ID.
46

47 In addition, if Drni_Gateway_Conversation and Ipp_Other_Gateway_Conversation are in
48 disagreement for any Gateway Conversation ID:

49 It sets DRF_Home_Oper_DRCP_State.Gateway_Sync to FALSE, and;

50 NTTDRCPDU to TRUE.
51

Otherwise:

52 DRF_Home_Oper_DRCP_State.Gateway_Sync and NTTDRCPDU are left unchanged.
53
54

1 Ipp_Gateway_Conversation_Direction is initialized to FALSE and recomputed whenever any
2 of its contributing variables changes.

3 updateIPPPortConversationPasses

4 This function computes a value for Ipp_Port_Conversation_Passes as follows:

5 For each Port Conversation ID, the value is TRUE (ID passes) if and only if:

- 6 a) the variables Drni_Port_Conversation and Ipp_Portal_System_State[] indicate that the
7 target Portal System for this Port Conversation ID lies behind this IPP, and
- 8 b) Drni_Port_Conversation and Ipp_Other_Port_Conversation_Portal_System are in
9 agreement as to which Portal System should get this Port Conversation ID.

10
11 In addition if Drni_Port_Conversation and Ipp_Other_Port_Conversation_Portal_System are in
12 disagreement for any Port Conversation ID:

13 It sets DRF_Home_Oper_DRCP_State.Port_Sync to FALSE, and;
14 NTTDRCPDU to TRUE.

15 Otherwise:

16 DRF_Home_Oper_DRCP_State.Port_Sync and NTTDRCPDU are left unchanged.

17
18 Ipp_Port_Conversation_Passes is initialized to FALSE and recomputed whenever any of its
19 contributing variables changes.

20 updateKey

21 This function updates the operational Aggregator Key, DRF_Home_Oper_Aggregator_Key, as
22 follows:

23 If Partner_DWC == TRUE then:

24 DRF_Home_Oper_Aggregator_Key is set to the value of the least significant 14 bits of
25 DRF_Home_Admin_Aggregator_Key, with the most significant two bits set to 01 so that
26 the same operational Key will be used by all the Portal Systems in the Portal;

27 Otherwise, if PSI == TRUE and Home Gateway == FALSE then:

28 DRF_Home_Oper_Aggregator_Key is set to the value of the least significant 14 bits of
29 DRF_Home_Admin_Aggregator_Key, with the most significant two bits set to 00 so that
30 the associated links cannot be part of the LAG that is set between the Partner Portals
31 (6.4.14.1).

32 Otherwise

33 DRF_Home_Oper_Aggregator_Key and all the operational keys of the Aggregation Ports
34 on this system are set to the lowest numerical non zero value of the set comprising the
35 values of the DRF_Home_Admin_Aggregator_Key, the
36 DRF_Neighbor_Admin_Aggregator_Key and the
37 DRF_Other_Neighbor_Admin_Aggregator_Key, on each IPP.

38 updateNTT

39 This function sets NTTDRCPDU to TRUE,
40 if any of:

- 41 DRF_Home_Oper_DRCP_State.Gateway_Sync, or;
 - 42 DRF_Home_Oper_DRCP_State.Port_Sync, or;
 - 43 DRF_Neighbor_Oper_DRCP_State.Gateway_Sync, or;
 - 44 DRF_Neighbor_Oper_DRCP_State.Port_Sync;
- 45 is FALSE.

46 updatePortalState

47 On all operations associated with this function, information provided by the
48 DRF_Other_Neighbor_State on an IPP is applicable only if Drni_Neighbor_ONN on the same
49 IPP is FALSE.

50
51 This function updates the Drni_Portal_System_State[] as follows:

52 The information for this Portal System, DRF_Home_State, indexed by the Portal System
53 Number, is included in Drni_Portal_System_State[]. For each of the other Portal Systems in the
54 Portal:

1 If any of the other Portal System's state information is available from two IPPs in this Portal
2 System, then:

3 For that Portal System, only the Portal System state information provided by the
4 DRF_Neighbor_State on the IPP having the other Portal System as a Neighbor Portal
5 System will be included in Drni_Portal_System_State[], indexed by the Portal System
6 Number.

7 Otherwise if a Portal System's state information is available only from a single IPP on this
8 Portal System, then:

9 that Portal System's state information, indexed by the associated Portal System Number
10 will be included in the Drni_Portal_System_State[] irrespectively of whether that
11 information is being provided by the DRF_Neighbor_State or the
12 DRF_Other_Neighbor_State on this IPP. If information for a Portal System is available
13 only from the DRF_Other_Neighbor_State on this IPP then ONN is set to TRUE on this
14 IPP.

15 Every Portal System included in the Portal Topology for which Portal System state information
16 is not available from any of the IPPs, has its associated Portal System state information in
17 Drni_Portal_System_State[] set to NULL (the Gateway is set to FALSE and the list of the
18 operational Aggregation Ports on the Portal System is emptied). For a Neighbor Portal System
19 for which the DRF_Neighbor_State is NULL, Drni_Neighbor_State[] is set equal to
20 Drni_Portal_System_State[].

21
22 This function updates also the Ipp_Portal_System_State[] for each IPP on this Portal System as
23 follows:

24 If any other Portal System's state information is available from two IPPs, then:

25 for every IPP on the Portal System, only the Portal System state information provided by
26 the DRF_Neighbor_State on that IPP will be included in the associated
27 Ipp_Portal_System_State[], indexed by the associated Portal System Number;

28 Otherwise;

29 the DRF_Neighbor_State on an IPP, indexed by the associated Portal System Number,
30 will be included as a first state in the corresponding Ipp_Portal_System_State[] and any
31 other additional state associated with another Portal System reported on the received
32 DRCPDU on this IPP, indexed by the associated Portal System Number, will be included
33 as the second state in the Ipp_Portal_System_State[].

34 Similarly to the Drni_Portal_System_State[], every Portal System included in the Portal
35 Topology for which Portal System state information is not available from any of the IPPs, has
36 its associated Portal System state information Ipp_Portal_System_State[] set to NULL (the
37 Gateway is set to FALSE, the Gateway Vector is set to NULL, and the list of the operational
38 Aggregation Ports on the Portal System is emptied).

39 updatePortalSystemGatewayConversation

40 This function sets Drni_Portal_System_Gateway_Conversation as follows:

41 If Differ_Gateway_Digest == TRUE and Drni_Three_System_Portal == 0;

42 Drni_Portal_System_Gateway_Conversation = Drni_Neighbor_Gateway_Conversation;

43 Otherwise;

44 This function sets the Drni_Portal_System_Gateway_Conversation to the result of the
45 logical AND operation between, the Boolean vector constructed from the
46 Drni_Gateway_Conversation, by setting to FALSE all the indexed Gateway
47 Conversation ID entries that are associated with other Portal Systems in the Portal, and the
48 Boolean vectors constructed from all IPPs Ipp_Other_Gateway_Conversation, by setting
49 to FALSE all the indexed Gateway Conversation ID entries that are associated with other
50 Portal Systems in the Portal.

51 updatePortalSystemPortConversation

52 This function sets Drni_Portal_System_Port_Conversation as follows:

53 If Differ_Port_Digest == TRUE and Drni_Three_System_Portal == 0;

54 Drni_Portal_System_Port_Conversation = Drni_Neighbor_Port_Conversation;

1 Otherwise;

2 This function sets the `Drni_Portal_System_Port_Conversation` to the result of the logical
3 AND operation between, the Boolean vector constructed from the
4 `Drni_Port_Conversation`, by setting to FALSE all the indexed Port Conversation ID
5 entries that are associated with other Portal Systems in the Portal, and the Boolean vector
6 constructed from the `Ipp_Other_Port_Conversation_Portal_System`, by setting to FALSE
7 all the indexed Port Conversation ID entries that are associated with other Portal Systems
8 in the Portal.
9

10 **9.4.12 Timers**

11 `DRCP_current_while_timer`

12 This timer is used to detect whether received protocol information has expired. If
13 `DRF_Home_Oper_DRCP_State.DRCP_Timeout` is set to Short Timeout, the timer is started
14 with the value `Short_Timeout_Time`. Otherwise, it is started with the value
15 `Long_Timeout_Time` (see 9.4.6).

16 `DRCP_periodic_timer (time_value)`

17 This timer is used to generate periodic transmissions. It is started using the value
18 `Slow_Periodic_Time` or `Fast_Periodic_Time` (see 9.4.6), as specified in the Periodic
19 Transmission state machine.
20
21

22 **9.4.13 Messages**

23 `DRCPCtrlMuxN:M_UNITDATA.indication(DRCPDU)`

24 This message is generated by the DRCP Control Parser/Multiplexer as a result of the reception
25 of a DRCPDU, formatted as defined in 9.4.3.2.
26
27

28 **9.4.14 DRCPDU Receive machine**

29 The DRCPDU Receive machine shall implement the function specified in Figure 9-23 with its associated
30 parameters (9.4.6 through 9.4.13). There is one DRCPDU Receive machine per IPP in a Portal System.
31
32

33 On receipt of a DRCPDU, the state machine enters the `PORTAL_CHECK` state. The `recordPortalValues`
34 function checks if the DRCPDU is associated with this Portal. If not, the state machine will enter the
35 `REPORT_TO_MANAGEMENT` state and the received DRCPDU will be reported to the management
36 system. While in the `REPORT_TO_MANAGEMENT` state, the System will exit to the `PORTAL_CHECK`
37 state if a new DRCPDU is received or to the `EXPIRED` state if the `DRCP_current_while_timer` expires. If
38 the `recordPortalValues` identifies the received DRCPDU as associated with this Portal, it will enter the
39 `COMPATIBILITY_CHECK` state to be checked by the `recordPortalConfValues` function. This compares the
40 administratively configured values that are associated with this portal to the received information and if they
41 differ the System will enter the `REPORT_TO_MANAGEMENT` state and the mis-configured DRCPDU
42 will be reported to the management system. If the Portal System continues to receive DRCPDUs that do not
43 match the administratively configured expectations for a period longer than twice the Short Timeout the
44 state machine will transit to the `DEFAULTED` state and the current operational parameters for the Portal
45 System(s) on this IPP will be overwritten with administratively configured values and a Portal System
46 update will be triggered.
47

48 If the received DRCPDU is configured according to the expected values for this Portal the DRCPDU
49 Receive machine enters the `CURRENT` state.
50

51 The `recordNeighborState` function records the Neighbor Portal System's State information contained in the
52 DRCPDU in the Neighbor Portal System's State operational variables and compares it with its own (Home)
53 Portal System state variable. If they differ, triggers are set to notify the Neighbor Portal System but also local
54

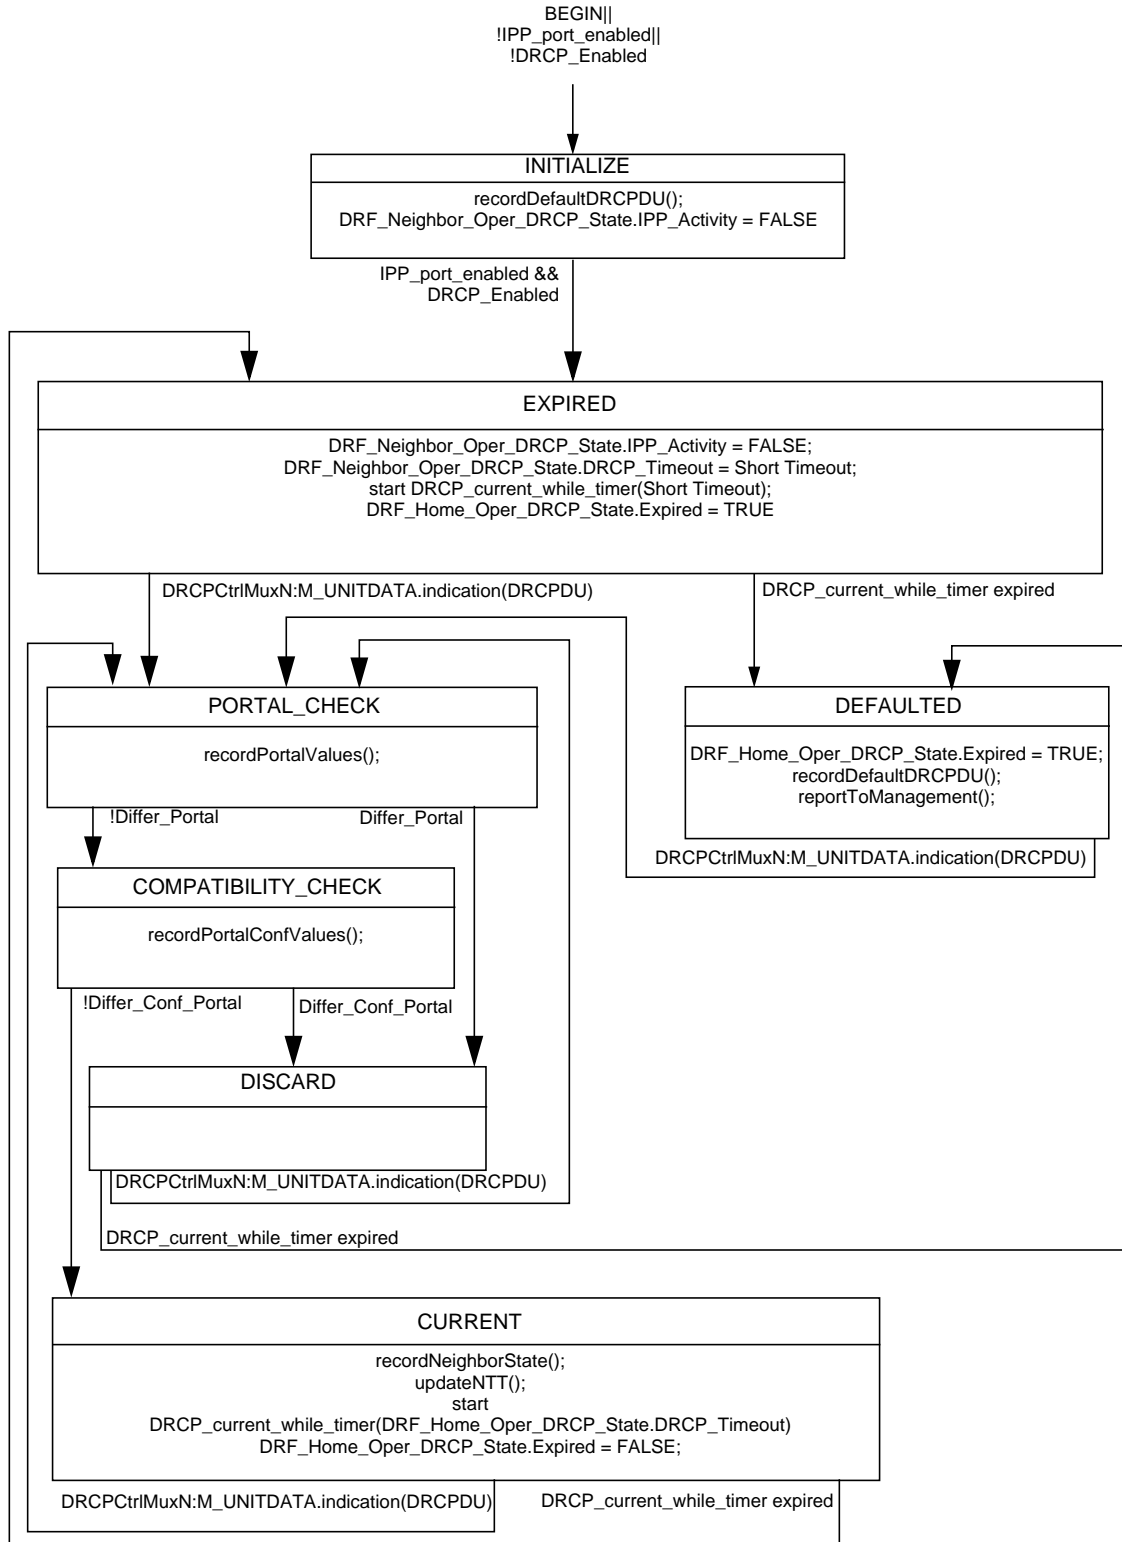


Figure 9-23—DRCPDU Receive machine state diagram

event variables are set to trigger updates on the local Portal System machine (PS—9.4.16), the DRNI Gateway and Aggregator machines (DGA—9.4.17), and the DRNI IPP machines (IPP—9.4.18).

In the process of executing the `recordPortalValues`, `recordPortalConfValues`, and `recordNeighborState`, functions, a DRCPDU Receive machine compliant to this standard shall not validate the Version Number, TLV_type, or Reserved fields in received DRCPDUs. The same actions are taken regardless of the values received in these fields. A DRCPDU Receive machine may validate the `Portal_Information_Length`, `Portal_Configuration_Information_Length`, `DRCP_State_Length`, or `Terminator_Length` fields. These behaviors, together with the constraint on future protocol enhancements, are discussed in 9.4.3.2.

NOTE—The rules expressed above allow Version 1 implementations to be compatible with future revisions of the protocol.

The `updateNTT` function is used to determine whether further protocol transmissions are required; `NTTDRCPDU` is set to TRUE if the Neighbor Portal System's view of the Home's operational Portal System state variable is not up to date. Then the `current_while` timer is started. The value used to start the timer is either `Short_Timeout_Time` or `Long_Timeout_Time`, depending upon the Portal System's operational value of `DRCP_Timeout`.

If no DRCPDU is received before the `DRCP_current_while_timer` expires, the state machine transits to the EXPIRED state. The `DRF_Neighbor_Oper_DRCP_State.IPP_Activity` is set to FALSE, the current operational value of the Neighbor Portal System's `DRCP_Timeout` variable is set to Short Timeout, and the `current_while` timer is started with a value of `Short_Timeout_Time`. This is a transient state; the `DRCP_Timeout` settings allow the Home Portal System to transmit DRCPDUs rapidly in an attempt to re-establish communication with the Neighbor Portal System.

If no DRCPDU is received before the `current_while` timer expires again, the state machine transits to the DEFAULTED state. The `recordDefaultDRCPDU` function overwrites the current operational parameters for the Neighbor Portal System with administratively configured values and triggers further updates to re-evaluate the Gateway and Port Conversation ID assignments and update its Aggregator's operational key value. The `reportToManagement` function is invoked in order to notify the management system for potential further actions.

If the IPP becomes inoperable, the state machine enters the INITIALIZE state. The `recordDefaultDRCPDU` function causes the administrative values of the Neighbor Portal System to be used as the current operational values, and the `DRF_Neighbor_Oper_DRCP_State.IPP_Activity` is set to FALSE. These actions force the PS machine to detach the Neighbor Portal System (and any Neighbor Portal System beyond it) from the Portal and set `GatewayConversationUpdate` and `PortConversationUpdate` to TRUE, which triggers a recomputation of the Gateway and Port Conversation ID filters.

9.4.15 DRCP Periodic Transmission machine

The DRCP Periodic Transmission machine shall implement the function specified in Figure 9-24 with its associated parameters (9.4.6 through 9.4.13). There is one DRCP Periodic Transmission machine per IPP in a Portal System.

The DRCP Periodic Transmission machine establishes the desire of the Home and the Neighbor Portal Systems to exchange periodic DRCPDUs on an IPP in order to maintain a Portal, and establishes how often those periodic transmissions should occur. Periodic transmissions will take place if either participant so wishes. Transmissions occur at a rate determined by the Neighbor Portal System; this rate is linked to the speed at which the Neighbor Portal System will time out received information.

The state machine has four states. They are as follows:

- a) *NO_PERIODIC*. While in this state, periodic transmissions are disabled.
- b) *FAST_PERIODIC*. While in this state, periodic transmissions are enabled at a fast transmission rate.

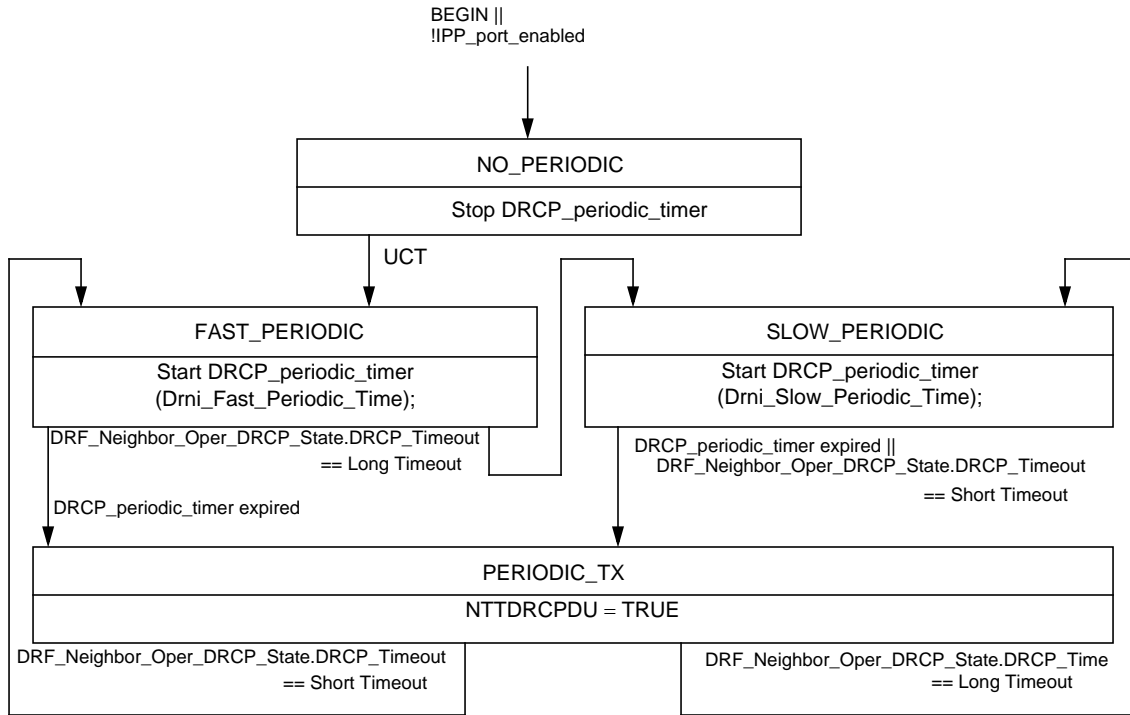


Figure 9-24—DRCP Periodic Transmission machine state diagram

- c) *SLOW_PERIODIC*. While in this state, periodic transmissions are enabled at a slow transmission rate.
- d) *PERIODIC_TX*. This is a transitory state entered on DRCP_periodic_timer expiry, that asserts NTTDRCPDU and then exits to FAST_PERIODIC or SLOW_PERIODIC depending upon the Neighbor Portal System’s DRCP_Timeout setting.

If periodic transmissions are enabled, the rate at which they take place is determined by the value of the DRF_Neighbor_Oper_DRCP_State.DRCP_Timeout variable. If this variable is set to Short Timeout, then the value fast_periodic_time is used to determine the time interval between periodic transmissions. Otherwise, slow_periodic_time is used to determine the time interval.

9.4.16 Portal System machine

The Portal System machine shall implement the function specified in Figure 9-25 with its associated parameters (9.4.6 through 9.4.13). There is one Portal System machine for each Portal System.

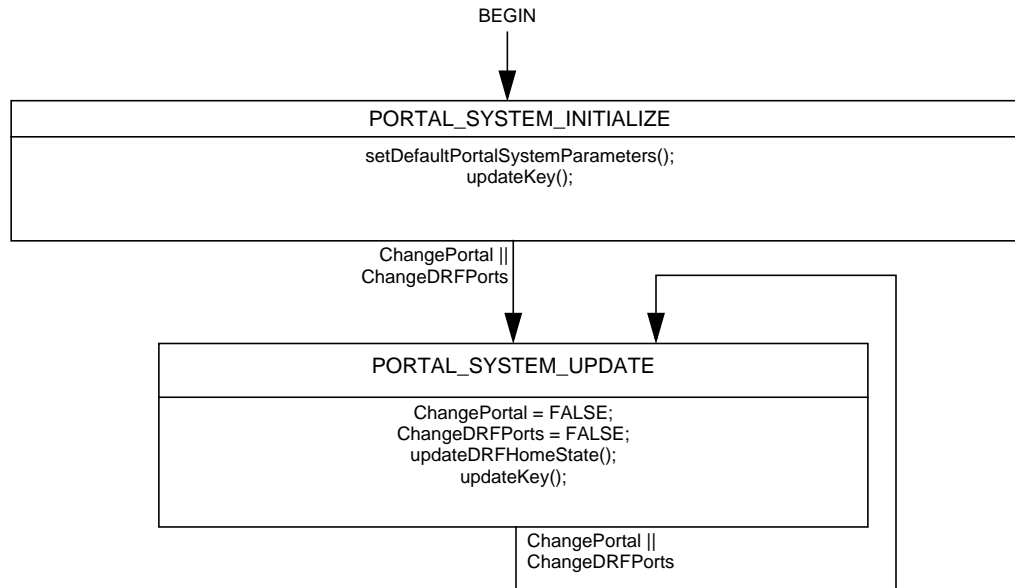


Figure 9-25—Portal System machine state diagram

The Portal System machine is responsible for updating the states of all Portal Systems in this Portal, based on information available to this Portal System.

On initialization the Portal System's variables are set to their default values for this Portal as configured by their administrative settings. In particular the default operational states of all Gateways, Aggregation Ports, and IPPs, in the Portal are set to FALSE. In addition based on those default values, the Operational Key to be used by the associated Aggregator is calculated to be the Administrative Key value assigned to this Portal System.

Any local change on the operational state of the Portal System's Gateway, or to the distribution state of any of the attached Aggregation Ports as reported by the associated Aggregator, or any change in the operational state of the Neighbor Portal Systems, as reported by the RX state machine, triggers a transition to the **PORTAL_SYSTEM_UPDATE** state. This causes the function `updateDRFHomeState` to re-evaluate the variable providing the Portal System's own state (`DRF_Home_State`) based on the updated local information on the operational state of the Gateway and all the Aggregation Ports on the Portal System's Aggregator. Any change in the operational state of the Portal System's Gateway is reflected to the `GatewayConversationUpdate` which is used to trigger state transitions in the ports' state machines (DGA (9.4.17) and IPP (9.4.18)). Similarly, any change in the operational state of the Aggregation Ports associated with this Portal System's Aggregator Port is reflected to the `PortConversationUpdate` which is used to trigger state transitions in the same state machines. Finally, the `updateKey` function updates the operational Key, to be used by the Portal System's Aggregator, by selecting the lowest numerical non zero value of the set comprising the values of the administrative Keys of all active Portal Systems in the Portal, if Conversation-sensitive frame collection and distribution (6.6) is not supported by the two Partner Systems, or by setting the operational Key of all Partner Systems in the Portal to same value if not.

The state machine returns to the **PORTAL_SYSTEM_UPDATE** state whenever the operational state of any of the DR Function's ports changes.

9.4.17 DRNI Gateway and Aggregator machines

The DRNI Gateway and Aggregator machines shall implement the function specified in Figure 9-26 with their associated parameters (9.4.6 through 9.4.13). There are two DRNI Gateway and Aggregator machines on a Portal System. Each one is associated with a Conversation ID type: there is one for Gateway Conversation IDs and one for Port Conversation IDs.

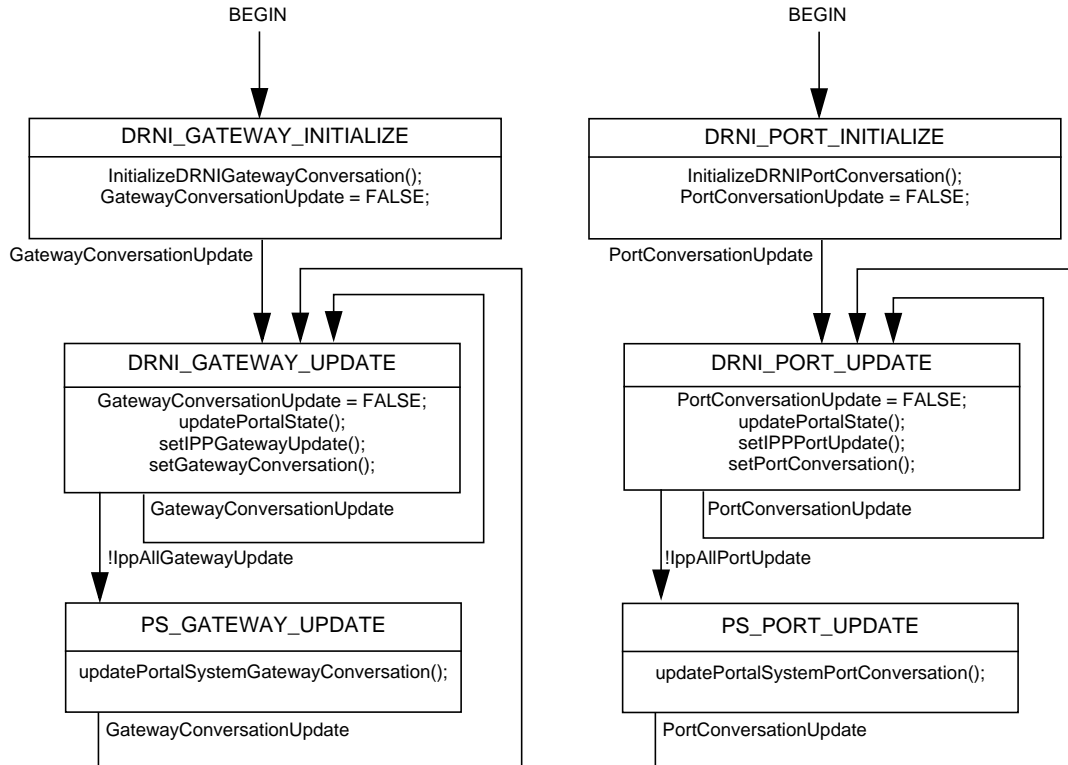


Figure 9-26—DRNI Gateway and Aggregator machines state diagrams

These state machines are responsible for configuring the Gateway Conversation IDs and the Port Conversation IDs that are allowed to pass through this DR Function's Gateway and Aggregator based on the agreed priority rules (7.3.1.1.35, 7.4.1.1.10) and the operation of DRCP.

On a trigger from the PS state machine (PS—9.4.16) or the DRX state machine (DRX—9.4.14), declaring that a Gateway's operational state has changed, the state machine enters the DRNI_GATEWAY_UPDATE state. This causes the triggering parameter (GatewayConversationUpdate) to be reset to FALSE while the function updatePortalState will update the variable providing the states of all Portal Systems (Drni_Portal_System_State[]) by combining the updated DRF_Home_State with information from the operational state of the ports on other Portal Systems as reported by the received DRCPDUs on the Portal System's IPPs and recorded by the DRX state machine (DRX—9.4.14). Subsequently, the function setIPPGatewayUpdate sets IppGatewayUpdate on every IPP on the Portal System to TRUE to trigger further updates on the IPP state machines (IPP—9.4.18) and the setGatewayConversation function is invoked to identify the Portal System that is responsible for each Gateway Conversation ID based on the agreed selection priorities and the Gateways operational state as known by this Portal System (based on the local Gateway's operational state and the Neighbor Portal Systems' declared operational state of their own Gateways carried by the latest DRCPDU received from those Neighbor Portal Systems). Finally the, Gateway Conversation ID indexed, Boolean vector will be calculated based on the agreement between this

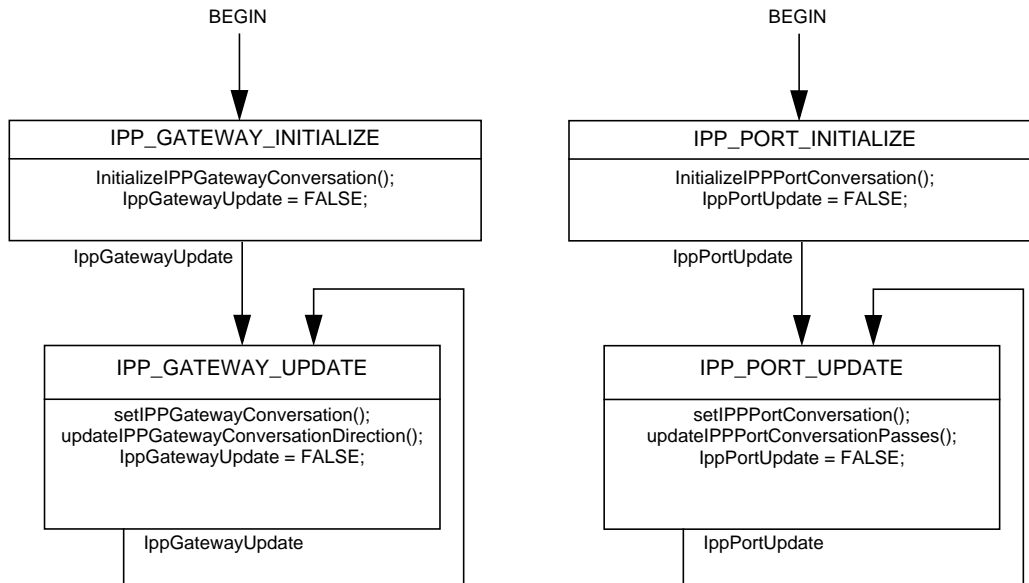
1 Portal System's view on the Gateway Conversation IDs that are allowed to pass through the Portal System's
2 Gateway and all the Neighbors' view on the Gateway Conversation IDs that are allowed to pass through this
3 Portal System's Gateway [as declared through their DRCPDUs and recorded by this Portal System's DRX
4 state machine (DRX—9.4.14)]. This ensures that no Gateway Conversation ID is allowed to pass through
5 this Portal System's Gateway unless agreement between all Portal Systems is reached.
6

7 The state machine is initialized having all Gateway Conversation IDs discarded and transits to the
8 DRNI_GATEWAY_UPDATE state whenever the trigger GatewayConversationUpdate is set.
9

10 The Port Conversation ID indexed Boolean vector is set through a similar state machine operation the only
11 difference being the priority selection rules are based on the agreed Port Conversation IDs and the Port
12 Algorithm instead of the agreed Gateway Conversation IDs and the Gateway Algorithm.
13

14 9.4.18 DRNI IPP machines

15
16 The DRNI IPP machines shall implement the function specified in Figure 9-27 with its associated
17 parameters (9.4.6 through 9.4.13). There is one DRNI IPP machine per each Conversation ID type per IPP in
18 a Portal System.
19



20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38 **Figure 9-27—DRNI IPP machine state diagrams**

39
40
41 These state machines are responsible for configuring the Gateway Conversation IDs and the Port
42 Conversation IDs that are allowed to pass through this Neighbor Portal System's IPPs based on the agreed
43 priority rules (7.3.1.1.35, 7.4.1.1.10) and the operation of DRCP.
44

45 On a trigger from the DRX state machine (DRX—9.4.14), declaring that IppGatewayUpdate is set to TRUE,
46 the state machine enters the IPP_GATEWAY_UPDATE state. Then the setIPPGatewayConversation
47 function will identify the Portal System that is responsible for each Gateway Conversation ID based on the
48 agreed selection priorities and the Gateways operational states as declared by the Neighbor Portal System on
49 this IPP (based on the Neighbor Portal System's Gateway operational state and the Neighbor Portal System's
50 declared operational state on their view on other Gateways in Portal, carried by the latest DRCPDU received
51 from the Neighbor Portal System on this IPP). Subsequently, Gateway Conversation ID indexed, Boolean
52 vector will be calculated based on the agreement between this Portal System's view on the Gateway
53 Conversation IDs that are allowed to pass through the Portal System's IPP and the IPP Neighbor Portal
54

System's view on the Gateway Conversation IDs that are allowed to pass through the same IPP [as declared through their DRCPDUs and recorded by this Portal System's DRX state machine (DRX—9.4.14)]. This ensures that no Gateway Conversation ID is allowed to pass through this IPP unless agreement between this Portal System and its Neighbor Portal System is reached. Finally, IppGatewayUpdate is reset to FALSE.

The state machine is initialized having all Gateway Conversation IDs discarded and transits to the IPP_GATEWAY_UPDATE state whenever the trigger GatewayConversationUpdate is set.

The Port Conversation ID indexed Boolean vector is set through a similar state machine operation, the only difference being that the priority selection rules are based on the agreed Port Conversation IDs and the Port Algorithm, instead of the agreed Gateway Conversation IDs and the Gateway Algorithm.

9.4.19 DRCPDU Transmit machine

When the DRCPDU Transmit machine creates a DRCPDU for transmission, it shall fill in the following fields with the corresponding operational values for this IPP:

- a) Aggregator ID and Priority.
- b) Portal Address and Priority.
- c) Portal System Number.
- d) Topology State.
- e) Operational Aggregator Key.
- f) Port algorithm.
- g) Gateway algorithm.
- h) Port Digest.
- i) Gateway Digest.
- j) If GatewayConversationTransmit is TRUE and if;
 - Drni_Three_System_Portal == 0;
 - The 2P Gateway Conversation Vector TLV ~~with its TLV length set to 512 Octets and its Gateway_Conversation field set to Drni_Portal_System_Gateway_Conversation~~ is prepared for DRCPDU transmission;
 - Otherwise if Drni_Three_System_Portal == 1;
 - The pair of 3P Gateway Conversation Vector-1 TLV and 3P Gateway Conversation Vector-2 TLV ~~with its TLV length set to 1024 Octets and its Gateway_Conversation field set to Drni_Gateway_Conversation~~ is are prepared for DRCPDU transmission;
- If PortConversationTransmit is TRUE and if;
 - Drni_Three_System_Portal == 0;
 - The 2P Port Conversation Vector TLV ~~with its TLV length set to 512 Octets and its Port_Conversation field set to Drni_Portal_System_Port_Conversation~~ is prepared for DRCPDU transmission;
 - Otherwise if Drni_Three_System_Portal == 1;
 - The 3P Port Conversation Vector-1 TLV accompanied by the 3P Port Conversation Vector-2 TLV ~~with its TLV length set to 1024 Octets and its Port_Conversation field set to Drni_Port_Conversation~~ is are prepared for DRCPDU transmission;
- If both GatewayConversationTransmit and PortConversationTransmit are TRUE and if;
 - Drni_Common_Methods is TRUE then;
 - either of the 2P Gateway Conversation Vector TLVs (the 2P Gateway Conversation Vector TLV, when Drni_Three_System_Portal == 0, or the pair of 3P Gateway Conversation Vector-1 TLV and 3P Gateway Conversation Vector-2 TLV, when Drni_Three_System_Portal == 1) or the 2P Port Conversation Vector TLVs (the 2P Port Conversation Vector TLV, when Drni_Three_System_Portal == 0, or the pair of 3P Port Conversation Vector-1 TLV and 3P Port Conversation Vector-2 TLV, when Drni_Three_System_Portal == 1) is are sufficient to be prepared for DRCPDU

transmission as each of the associated Conversation Vector TLVs set is applicable to both Gateway and Port distributions;
Otherwise if `Drni_Common_Methods` is FALSE and `Drni_Three_System_Portal == 1`;
Two separate DRCPDUs, one including the pair of 3P Gateway Conversation Vector-1 TLV and 3P Gateway Conversation Vector-2 TLV and an other one including the pair of 3P Port Conversation Vector-1 TLV and 3P Port Conversation Vector-2 TLV will be prepared for transmission having all other TLVs the same.

If present, the Conversation Vector TLVs should be inserted between the Conversation Vector Indicator TLV and the DRCP State TLV.

- k) DRCP State.
- l) The operational Aggregation Ports, the Administrative Aggregator Key and the operational Partner Aggregator Key of the Home Portal System and any other Portal System in the `Drni_Portal_System_State[]` (9.4.8).
- m) The operational Gateway Sequence numbers for every Portal System in the `Drni_Portal_System_State[]` and, if `HomeGatewayVectorTransmit` and/or `OtherGatewayVectorTransmit` are set to TRUE and the potential presence of the Conversation Vector TLVs does not result in a DRCPDU that has a length that is larger than the maximum allowed by the access method supporting that IPP, the associated Home Gateway Vector and/or Other Gateway Vector respectively.

In addition if the system is configured to use one of the Network/IPL shared methods specified in 9.3.2.1, 9.3.2.2, or 9.3.2.3 by configuring a non-NULL value in `aDrniEncapsulationMethod` additional Network/IPL sharing TLVs will need to be attached to the main DRCPDU, carrying the appropriate operational values as specified in 9.4.3.4.1 and 9.4.3.4.2

When the Periodic machine is in the NO_PERIODIC state, the DRCPDU Transmit machine shall

- Not transmit any DRCPDUs, and
- Set the value of `NTTDRCPDU` to FALSE.

When the `DRCP_Enabled` variable is TRUE and the `NTTDRCPDU` (9.4.9) variable is TRUE, the Transmit machine shall ensure that a properly formatted DRCPDU (9.4.3.2) is transmitted [i.e., issue a `DRCPCtrlMuxN:M_UNITDATA.Request(DRCPDU)` service primitive]. The `NTTDRCPDU` variable shall be set to FALSE when the Transmit machine has transmitted a DRCPDU containing all the required fields as specified in this clause.

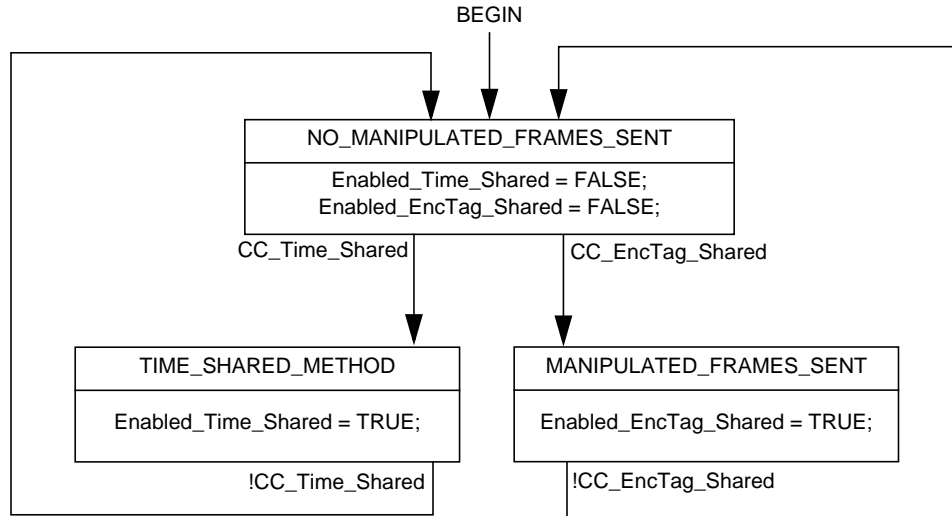
When the `DRCP_Enabled` variable is FALSE, the Transmit machine shall not transmit any DRCPDUs and shall set the value of `NTTDRCPDU` to FALSE.

9.4.20 Network/IPL sharing machine

The Network/IPL sharing machine shall implement the functions specified in Figure 9-28 with its associated parameters (9.4.6 through 9.4.13). There is one Network/IPL sharing machine per IPP in a Portal System for the supported Network/IPL sharing method. This machine is only required when the Network/IPL shared methods, Network / IPL sharing by time (9.3.2.1), Network / IPL sharing by tag (9.3.2.2), or Network / IPL sharing by encapsulation (9.3.2.3) is implemented.

The Network/IPL sharing machine enables transmission and manipulation of the frames sent on the shared Network/IPL link only if the DRCPDUs received on the same port report the same Network/IPL sharing configuration by the Neighbor Portal System.

The state machine has three states. They are as follows:



19 **Figure 9-28—Network/IPL sharing machine state diagram**

- 20
21
22
23
24
25
26
27
- NO_MANIPULATED_FRAMES_SENT. While in this state, the IPL can only be supported by a physical or Aggregation Link.
 - TIME_SHARED_METHOD. While in this state, the Network / IPL sharing by time methods specified in 9.3.2.1 are enabled.
 - MANIPULATED_FRAMES_SENT. While in this state, the tag manipulation methods of Network / IPL sharing by tag or Network / IPL sharing by encapsulation, as dictated by the Network / IPL sharing method selected the aDrniEncapsulationMethod (7.4.1.1.17), are enabled.

28
29
30
31
32
33
34
35
36

The System is initialized in the NO_MANIPULATED_FRAMES_SENT and IPL frames are sent on the dedicated physical link. If the Home Portal System is configured for Network / IPL sharing by time mode of operation, indicated by a value of 1 in aDrniEncapsulationMethod (7.4.1.1.17), the system will transit to TIME_SHARED_METHOD if the DRX state machine (DRX—9.4.14) sets CC_Time_Shared to TRUE (indicating that the Neighbor Portal System on this IPP has also been configured for the Network / IPL sharing by time mode of operation). The System remains in the TIME_SHARED_METHOD state until a received DRCPDU sets CC_Time_Shared to FALSE, which triggers a state transition to the NO_MANIPULATED_FRAMES_SENT state and IPL frames are sent on the dedicated physical link.

37
38
39
40
41
42
43
44
45

Similarly, if the Home Portal System is configured for Network / IPL sharing by tag or Network / IPL sharing by encapsulation mode of operation, as indicated by the value in aDrniEncapsulationMethod (7.4.1.1.17), the system will transit to MANIPULATED_FRAMES_SENT if the DRX state machine (DRX—9.4.14) sets CC_EncTag_Shared to TRUE (indicating that the Neighbor Portal System on this IPP has also been configured for the Network / IPL sharing by tag or Network / IPL sharing by encapsulation mode of operation respectively). The System remains in the MANIPULATED_FRAMES_SENT state until a received DRCPDU sets CC_EncTag_Shared to FALSE, which triggers a state transition to the NO_MANIPULATED_FRAMES_SENT state and IPL frames are sent on the dedicated physical link.

Annex A

(normative)

Protocol Implementation Conformance Statement (PICS) proforma

A.1 Introduction¹

The supplier of an implementation that is claimed to conform to Clause 6, Link Aggregation, shall complete the following protocol implementation conformance statement (PICS) proforma.

A completed PICS proforma is the PICS for the implementation in question. The PICS is a statement of which capabilities and options of the protocol have been implemented. A PICS is included at the end of each clause as appropriate. The PICS can be used for a variety of purposes by various parties, including the following:

- a) As a checklist by the protocol implementer, to reduce the risk of failure to conform to the standard through oversight;
- b) As a detailed indication of the capabilities of the implementation, stated relative to the common basis for understanding provided by the standard PICS proforma, by the supplier and acquirer, or potential acquirer, of the implementation;
- c) As a basis for initially checking the possibility of interworking with another implementation by the user, or potential user, of the implementation (note that, while interworking can never be guaranteed, failure to interwork can often be predicted from incompatible PICS);
- d) As the basis for selecting appropriate tests against which to assess the claim for conformance of the implementation, by a protocol tester.

A.1.1 Abbreviations and special symbols

The following symbols are used in the PICS proforma:

M	mandatory field/function
!	negation
O	optional field/function
O.<n>	optional field/function, but at least one of the group of options labeled by the same numeral <n> is required
O/<n>	optional field/function, but one and only one of the group of options labeled by the same numeral <n> is required
X	prohibited field/function
<item>:	simple-predicate condition, dependent on the support marked for <item>
<item1>*<item2>:	AND-predicate condition, the requirement shall be met if both optional items are implemented

1. *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

A.1.2 Instructions for completing the PICS proforma

The first part of the PICS proforma, Implementation Identification and Protocol Summary, is to be completed as indicated with the information necessary to identify fully both the supplier and the implementation.

The main part of the PICS proforma is a fixed-format questionnaire divided into subclauses, each containing a group of items. Answers to the questionnaire items are to be provided in the right-most column, either by simply marking an answer to indicate a restricted choice (usually Yes, No, or Not Applicable), or by entering a value or a set or range of values. (Note that there are some items where two or more choices from a set of possible answers can apply; all relevant choices are to be marked.)

Each item is identified by an item reference in the first column; the second column contains the question to be answered; the third column contains the reference or references to the material that specifies the item in the main body of the standard; the sixth column contains values and/or comments pertaining to the question to be answered. The remaining columns record the status of the items—whether the support is mandatory, optional or conditional—and provide the space for the answers.

The supplier may also provide, or be required to provide, further information, categorized as either Additional Information or Exception Information. When present, each kind of further information is to be provided in a further subclause of items labeled A<*i*> or X<*i*>, respectively, for cross-referencing purposes, where <*i*> is any unambiguous identification for the item (e.g., simply a numeral); there are no other restrictions on its format or presentation.

A completed PICS proforma, including any Additional Information and Exception Information, is the protocol implementation conformance statement for the implementation in question.

Note that where an implementation is capable of being configured in more than one way, according to the items listed under Major Capabilities/Options, a single PICS may be able to describe all such configurations. However, the supplier has the choice of providing more than one PICS, each covering some subset of the implementation's configuration capabilities, if that would make presentation of the information easier and clearer.

A.1.3 Additional information

Items of Additional Information allow a supplier to provide further information intended to assist the interpretation of the PICS. It is not intended or expected that a large quantity will be supplied, and the PICS can be considered complete without any such information. Examples might be an outline of the ways in which a (single) implementation can be set up to operate in a variety of environments and configurations; or a brief rationale, based perhaps upon specific application needs, for the exclusion of features that, although optional, are nonetheless commonly present in implementations.

References to items of Additional Information may be entered next to any answer in the questionnaire, and may be included in items of Exception Information.

A.1.4 Exceptional information

It may occasionally happen that a supplier will wish to answer an item with mandatory or prohibited status (after any conditions have been applied) in a way that conflicts with the indicated requirement. No preprinted answer will be found in the Support column for this; instead, the supplier is required to write into the Support column an X<*i*> reference to an item of Exception Information, and to provide the appropriate rationale in the Exception item itself.

1 An implementation for which an Exception item is required in this way does not conform to this standard.

2
3 Note that a possible reason for the situation described above is that a defect in the standard has been
4 reported, a correction for which is expected to change the requirement not met by the implementation.
5

6 7 **A.1.5 Conditional items**

8
9 The PICS proforma contains a number of conditional items. These are items for which both the applicability
10 of the item itself, and its status if it does apply—mandatory, optional, or prohibited—are dependent upon
11 whether or not certain other items are supported.
12

13 Individual conditional items are indicated by a conditional symbol of the form “<item>:<s>” in the Status
14 column, where “<item>” is an item reference that appears in the first column of the table for some other
15 item, and “<s>” is a status symbol, M (Mandatory), O (Optional), or X (Not Applicable).
16

17 If the item referred to by the conditional symbol is marked as supported, then 1) the conditional item is
18 applicable, 2) its status is given by “<s>”, and 3) the support column is to be completed in the usual way.
19 Otherwise, the conditional item is not relevant and the Not Applicable (N/A) answer is to be marked.
20

21 Each item whose reference is used in a conditional symbol is indicated by an asterisk in the Item column.
22

23 24 **A.1.6 Identification**

25 26 **A.1.6.1 Implementation identification**

28 Supplier (Note 1)	
29 Contact point for queries about the PICS (Note 1)	
30 Implementation Name(s) and Version(s) (Notes 1 and 3)	
31 Other information necessary for full identification—e.g., 32 name(s) and version(s) of machines and/or operating 33 system names (Note 2)	
34 35 36 NOTE 1—Required for all implementations. 37 NOTE 2—May be completed as appropriate in meeting the requirements for the identification. 38 NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s 39 terminology (e.g., Type, Series, Model).	

40 41 42 **A.1.6.2 Protocol summary**

44 Identification of protocol specification	IEEE Std 802.1AX-2008, Clause 6, Link Aggregation.
45 Identification of amendments and corrigenda to the 46 PICS proforma that have been completed as part of 47 the PICS	

A.2 PICS proforma for Clause 6, Link Aggregation

A.2.1 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
	The items below are relevant to the ports of a System for which support is claimed				
LA	Is Link Aggregation supported as specified in Clause 6	item a) in 5.3, 6.2	Support the Link Aggregation Sublayer and conform to the state machines and procedures in 6.2	M	Yes []
LACP	Does the Link Aggregation Control Protocol	item b) in 5.3, 6.3, 6.4	Support the Link Aggregation Control Protocol and conform to the state machines and procedures in 6.3 and 6.4	M	Yes []
V1LA	Are Version 1 LACPDUs supported?	item c) in 5.3, 6.4.2	Transmit and receive version 1 LACPDUs in the formats specified in 6.4.2	M	Yes []
V2LA	Are Version 2 LACPDUs supported	item h) in 5.3.1, 6.4.2	Transmit and receive version 2 LACPDUs in the formats specified in 6.4.2	O CSCD3: M	Yes [] No []
MG	Is the Marker Generator/ Receiver supported?	item a) in 5.3.1, 6.2.5		O	Yes [] No []
MGT	Is Management supported?	item b) in 5.3.1, Clause 7	Support the management functionality for Link Aggregation as specified in Clause 7	O	Yes [] No []
MIB	Does the implementation support management operations using SMIPv2 MIB modules?	item c) in 5.3.1, Annex D	Support SMIPv2 MIB modules for the management of Link Aggregation capabilities	MGT: O	N/A [] Yes [] No []
AM	Is there Aggregation Port Debug Information package support?	item d) in 5.3.1, 7.3		MGT: O	N/A [] Yes [] No []
CM	Does the implementation support the Churn Detection machine?	item d) in 5.3.1, 6.4.17	Required if Aggregation Port Debug Information package supported	AM: M !AM: O	N/A [] Yes [] No []
PSFD	Is Per-service frame distribution supported?	item g) in 5.3.1, 8.2		O	Yes [] No []
CSCD	Is Conversation-sensitive frame collection and distribution supported?	item g) in 5.3.1, A.2.27		PSFD: M	N/A [] Yes []
DRNI	Is Distributed Resilient Network Interconnect supported?	5.4, A.2.30		O	Yes [] No []

A.2.2 LLDP Port connectivity

Item	Feature	Subclause	Value/Comment	Status	Support
LLDP	Does the System support LLDP?	item e) in 5.3.1, 6.1.3		O	Yes [] No []
LLDPP M	Does the System support LLDP Parser/Multiplexers on the Aggregation Ports?	item e) in 5.3.1, 6.1.3		LLDP: M	N/A [] Yes []

A.2.3 Protocol Parser/Multiplexer support

Item	Feature	Subclause	Value/Comment	Status	Support
PPM	Does the System support Protocol Parser/Multiplexer for a protocol that is supported by the System but not specified in this standard?	item f) in 5.3.1, 6.2.7		O	N/A [] Yes []

A.2.4 Frame Collector

Item	Feature	Subclause	Value/Comment	Status	Support
FC1	Frame Collector function	6.2.3	As specified in the state machine shown in Figure 6-3 and associated definitions in 6.2.3.1	M	Yes []
FC2	Frame Collector function—CollectorMaxDelay	6.2.3.1.4	Deliver or discard frames within CollectorMaxDelay	M	Yes []

A.2.5 Frame Distributor

Item	Feature	Subclause	Value/Comment	Status	Support
	Distribution algorithm ensures the following, when frames received by Frame Collector:	6.2.4			

Item	Feature	Subclause	Value/Comment	Status	Support
FD1	Frame misordering		None	M	Yes []
FD2	Frame duplication		None	M	Yes []
FD3	Frame Distributor function	6.2.4	Function as specified in the state machine shown in Figure 6-4 and associated definitions in 6.2.4.1	M	Yes []

A.2.6 Marker protocol

Item	Feature	Subclause	Value/Comment	Status	Support
MGR1	Marker Generator/Receiver	6.2.5		MG:M	N/A [] Yes []
MGR2	Marker Responder	item d) in 5.3, 6.2.6	Function specified in 6.5.4.2	!DRN1:M DRN1:O	N/A [] Yes [] No []

A.2.7 Aggregator Parser/Multiplexer

Item	Feature	Subclause	Value/Comment	Status	Support
APM1	Aggregator Multiplexer	6.2.8	Transparent pass-through of frames	M	Yes []
APM2	Aggregator Multiplexer	6.2.8	Discard of TX frames when Aggregation Port not Distributing	M	Yes []
APM3	Aggregator Parser	6.2.8	Function specified by state machine shown in Figure 6-6 and associated definitions in 6.2.8.1	M	Yes []
APM4	Aggregator Parser	6.2.8	Discard of RX frames when Aggregation Port not Collecting	M	Yes []

A.2.8 Control Parser/Multiplexer

Item	Feature	Subclause	Value/Comment	Status	Support
CPM1	Control Multiplexer	6.2.10	Transparent pass-through of frames	M	Yes []
CPM2	Control Parser	6.2.10	Function specified by state machine shown in Figure 6-5 and associated definitions in 6.2.10.1 and 6.2.9	M	Yes []

A.2.9 System identification

Item	Feature	Subclause	Value/Comment	Status	Support
SID1	Globally unique identifier	6.3.2	Globally administered individual MAC address plus System Priority	M	Yes []
SID2	MAC address chosen	6.3.2	MAC address associated with one of the Aggregation Ports	O	Yes [] No []

A.2.10 Aggregator identification

Item	Feature	Subclause	Value/Comment	Status	Support
AID1	Globally unique identifier	6.3.3	Globally administered individual MAC address	M	Yes []
AID2	Integer identifier	6.3.3	Uniquely identifies the Aggregator within the System	M	Yes []
*AID3	Unique identifier allocated	6.3.3	Unique identifier assigned to one of its bound Aggregation Ports	O	Yes [] No []
AID4			Unique identifier not assigned to any other Aggregator	!AID3 :M	N/A [] Yes []

A.2.11 Port identification

Item	Feature	Subclause	Value/Comment	Status	Support
PID1	Port Identifiers	6.3.4	Unique within a System; Port Number 0 not used for any Aggregation Port	M	Yes []

A.2.12 Capability identification

Item	Feature	Subclause	Value/Comment	Status	Support
CID1	Administrative and operational Key values associated with each Aggregation Port	6.3.5		M	Yes []
CID2	Administrative and operational Key values associated with each Aggregator	6.3.5		M	Yes []

A.2.13 Link Aggregation Group identification

Item	Feature	Subclause	Value/Comment	Status	Support
LAG1	LAG ID component values	6.3.6.1	Actor's values non-zero. Partner's admin values only zero for Individual Aggregation Ports	M	Yes []

A.2.14 Detaching a link from an Aggregator

Item	Feature	Subclause	Value/Comment	Status	Support
DLA1	Effect on conversation reallocated to a different link	6.3.14	Frame ordering preserved	M	Yes []

A.2.15 LACPDU structure

Item	Feature	Subclause	Value/Comment	Status	Support
LPS2	LACPDU structure	6.4.2.3	As shown in Figure 6-7 and as described	M	Yes []
LPS3	LACPDU structure	6.4.2	All Reserved octets ignored on receipt and transmitted as zero	M	Yes []

A.2.16 Version 2 LACPDU

Item	Feature	Subclause	Value/Comment	Status	Support
	If V2LA is not supported, mark N/A and ignore the remainder of this table				N/A []
V2LA1	Is the Long LACPDU machine supported?	6.4.18	As shown in Figure 6-25 and as described	V2LA:M	Yes []
V2LA2	Is the Port Algorithm TLV supported?	6.4.2.4.1	As shown in Figure 6-9 and as described	V2LA:M	Yes []
V2LA3	Is the Port Conversation ID Digest TLV supported?	6.4.2.4.2	As shown in Figure 6-10 and as described	V2LA:M	Yes []
V2LA4	Is the Port Conversation Mask TLVs supported?	6.4.2.4.3	As shown in Figure 6-11, Figure 6-12, Figure 6-13, Figure 6-14 and as described	V2LA:M	Yes []
V2LA5	Is the Port Conversation Service Mapping TLV supported?	6.4.2.4.4	As shown in Figure 6-16 and as described	V2LA:O V2LA AND PSFD2: M	Yes [] No []

A.2.17 State machine variables

Item	Feature	Subclause	Value/Comment	Status	Support
SMV1	Partner_Admin_Port_State	6.4.7	Collecting set to the same value as Synchronization	M	Yes []
SMV2	LACP_Enabled	6.4.8	TRUE for point-to-point links, otherwise FALSE	M	Yes []

A.2.18 Receive machine

Item	Feature	Subclause	Value/Comment	Status	Support
RM1	Receive machine	6.4.12	As defined in Figure 6-18 and associated parameters	M	Yes []
RM2	Validation of LACPDUs		No validation of Version Number, TLV_type, or Reserved fields	M	Yes []

A.2.19 Periodic Transmission machine

Item	Feature	Subclause	Value/Comment	Status	Support
PM1	Periodic Transmission machine	6.4.13	As defined in Figure 6-19 and associated parameters	M	Yes []

A.2.20 Selection Logic

Item	Feature	Subclause	Value/Comment	Status	Support
	Selection logic requirements	6.4.14.1			
SLM1	Aggregator support		At least one Aggregator per System	M	Yes []
SLM2	Aggregation Port Keys		Each Aggregation Port assigned an operational Key	M	Yes []
SLM3	Aggregator Keys		Each Aggregator assigned an operational Key	M	Yes []
SLM4	Aggregator Identifiers		Each Aggregator assigned an identifier	M	Yes []
SLM5	Aggregator selection		If same Key assignment as Aggregation Port	M	Yes []
SLM6	Aggregation Ports that are members of the same Link Aggregation Group		Aggregation Ports select same Aggregator	M	Yes []
SLM7	Pair of Aggregation Ports connected in loopback		Not select same Aggregator as each other	M	Yes []
SLM8	Aggregation Port required to be Individual		Not select same Aggregator as any other Aggregation Port	M	Yes []

Item	Feature	Subclause	Value/Comment	Status	Support
SLM9	Aggregation Port is Aggregate-able		Not select same Aggregator as any Individual Aggregation Port	M	Yes []
SLM10	Aggregation Port unable to select an Aggregator		Aggregation Port not attached to any Aggregator	M	Yes []
SLM11	Further aggregation constraints		Aggregation Ports may be selected as standby	O	Yes [] No []
SLM12	Selected variable		Set to SELECTED or STANDBY once Aggregator is determined	M	Yes []
SLM13	Port enabled		Only when selected and attached to an Aggregator	M	Yes []
SLM14	Recommended default operation of Selection Logic	6.4.14.2	Meets requirements of 6.4.14.2	O	Yes [] No []

A.2.21 Mux machine

Item	Feature	Subclause	Value/Comment	Status	Support
XM1	Mux machine	6.4.15	As defined in Figure 6-21 or Figure 6-22, and associated parameters	M	Yes []

A.2.22 Transmit machine

Item	Feature	Subclause	Value/Comment	Status	Support
	Transmitted in outgoing LACP-DUs	6.4.16			
TM1	Actor_Port and Actor_Port_Priority	6.4.16		M	Yes []
TM2	Actor_System and Actor_System_Priority	6.4.16		M	Yes []
TM3	Actor_Key	6.4.16		M	Yes []
TM4	Actor_State	6.4.16		M	Yes []
TM5	Partner_Port and Partner_Port_Priority	6.4.16		M	Yes []
TM6	Partner_System and Partner_System_Priority	6.4.16		M	Yes []

Item	Feature	Subclause	Value/Comment	Status	Support
TM7	Partner_Key	6.4.16		M	Yes []
TM8	Partner_State	6.4.16		M	Yes []
TM9	CollectorMaxDelay	6.4.16		M	Yes []
TM10	Action when Periodic machine is in the NO_PERIODIC state	6.4.16	Set NTT to FALSE, do not transmit	M	Yes []
TM11	Action when LACP_Enabled is TRUE, NTT is TRUE, and not rate limited	6.4.16	Properly formatted LACPDU transmitted	M	Yes []
TM12	Action when LACP_Enabled is TRUE and NTT is TRUE, when rate limit is in force	6.4.16	Transmission delayed until limit is no longer in force	M	Yes []
TM13	Action when LACPDU has been transmitted	6.4.16	Set NTT to FALSE	M	Yes []
TM14	Action when LACP_Enabled is FALSE	6.4.16	Set NTT to FALSE, do not transmit	M	Yes []

A.2.23 Churn Detection machines

Item	Feature	Subclause	Value/Comment	Status	Support
CM1	Churn Detection machines	6.4.17	As defined in Figure 6-23 and Figure 6-24	CM:M	N/A [] Yes []

A.2.24 Marker protocol

Item	Feature	Subclause	Value/Comment	Status	Support
FP1	Respond to all received Marker PDUs	item d) in 5.3, 6.5.1	As specified by 6.5.4	!DRN1:M DRN1:O	N/A [] Yes [] No []
FP2	Use of the Marker protocol	6.5.1	As specified by 6.5.4	O	Yes [] No []
FP3	MARKER.request service primitives request rate	6.5.4.1	Maximum of five during any one-second period	MG:M	N/A [] Yes []
FP6	Marker PDU structure	6.5.3.3	As shown in Figure 6-27 and as described	MG:M	N/A [] Yes []

Item	Feature	Subclause	Value/Comment	Status	Support
FP7	Marker Response PDU structure	item d) in 5.3, 6.5.3.3	As shown in Figure 6-27 and as described	!DRNI1:M DRNI1:O	N/A [] Yes [] No []
FP8	Marker Responder state machine	item d) in 5.3, 6.5.4.2	As specified in Figure 6-28 and 6.5.4.2.1 through 6.5.4.2.2	!DRNI1:M DRNI1:O	N/A [] Yes [] No []
FP9	Validation of Marker Request PDUs	item d) in 5.3, 6.5.4.2.2	Marker Responder shall not validate the Version Number, Pad, or Reserved fields	!DRNI1:M DRNI1:O	N/A [] Yes [] No []

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

A.2.25 Management

Item	Feature	Subclause	Value/Comment	Status	Support
	If MGT is not supported, mark N/A and ignore the remainder of this table				N/A []
MGT1	Is the Basic package supported?	Table 7–1, 7.3.1.1.1, 7.3.2.1.1	Support of the Managed objects marked as members of the Basic package in Table 7–1	MGT: M	N/A [] Yes []
MGT2	Is the Mandatory package supported?	Table 7–1, 7.3.1.1.2, 7.3.1.1.3, 7.3.1.1.4, 7.3.1.1.5, 7.3.1.1.6, 7.3.1.1.7, 7.3.1.1.8, 7.3.1.1.9, 7.3.1.1.10, 7.3.1.1.11, 7.3.1.1.12, 7.3.1.1.13, 7.3.1.1.14, 7.3.1.1.15, 7.3.1.1.16, 7.3.1.1.19, 7.3.1.1.20, 7.3.1.1.31, 7.3.1.2.1, 7.3.1.2.2, 7.3.1.1.32, 7.3.2.1.2, 7.3.2.1.3, 7.3.2.1.4, 7.3.2.1.5, 7.3.2.1.6, 7.3.2.1.7, 7.3.2.1.8, 7.3.2.1.9, 7.3.2.1.10, 7.3.2.1.11, 7.3.2.1.12, 7.3.2.1.13, 7.3.2.1.14, 7.3.2.1.15, 7.3.2.1.16, 7.3.2.1.17, 7.3.2.1.18, 7.3.2.1.19, 7.3.2.1.20, 7.3.2.1.21, 7.3.2.1.22, 7.3.2.1.23, 7.3.2.1.24, 7.3.2.2.1, 7.3.2.2.1	Support of the Managed objects marked as members of the Mandatory package in Table 7–1	MGT: M	N/A [] Yes []

Item	Feature	Subclause	Value/Comment	Status	Support
MGT3	Is the Recommended package supported?	Table 7–1, 7.3.1.1.17, 7.3.1.1.18, 7.3.1.1.25, 7.3.1.1.26, 7.3.1.1.27, 7.3.1.1.28, 7.3.1.1.29, 7.3.1.1.30	Support of the Managed objects marked as members of the Recommended package in Table 7–1	MGT: O	N/A [] Yes [] No []
MGT4	Is the Optional package supported?	Table 7–1, 7.3.1.1.21, 7.3.1.1.22, 7.3.1.1.23, 7.3.1.1.24	Support of the Managed objects marked as members of the Optional package in Table 7–1	MGT: O	N/A [] Yes [] No []
MGT5	Is the Aggregation Port Statistics package supported?	Table 7–1, 7.3.3.1.1, 7.3.3.1.2, 7.3.3.1.3, 7.3.3.1.4, 7.3.3.1.5, 7.3.3.1.6, 7.3.3.1.7, 7.3.3.1.8, 7.3.3.1.9	Support of the Managed objects marked as members of the Aggregation Port Statistics package in Table 7–1	MGT: O	N/A [] Yes [] No []
MGT6	Is the Aggregation Port Debug Information package supported?	Table 7–1, 7.3.4.1.1, 7.3.4.1.2, 7.3.4.1.3, 7.3.4.1.4, 7.3.4.1.5, 7.3.4.1.6, 7.3.4.1.7, 7.3.4.1.8, 7.3.4.1.9, 7.3.4.1.10, 7.3.4.1.11, 7.3.4.1.12, 7.3.4.1.13	Support of the Managed objects marked as members of the Aggregation Port Debug Information package in Table 7–1	MGT: O	N/A [] Yes [] No []
MGT7	Is the Per-Service Frame Distribution package supported?	Table 7–1, 7.3.1.1.33, 7.3.1.1.34, 7.3.1.1.35, 7.3.1.1.36, 7.3.1.1.37, 7.3.1.1.38, 7.3.1.1.39 7.3.2.1.25, 7.3.2.1.26, 7.3.2.1.27, 7.3.2.1.28, 7.3.2.1.29	Support of the Managed objects marked as members of the Per-Service Frame Distribution package in Table 7–1	MGT AND PSFD: M	N/A [] Yes []
MGT8	Are the CDS Churn Detection managed objects supported?	Table 7–1, 7.3.4.1.14, 7.3.4.1.15, 7.3.4.1.16, 7.3.4.1.17	These managed object are applicable only when Conversation-sensitive frame collection and distribution as specified in 6.6 is supported	MGT AND PSFD: O MGT6 AND PSFD: M	N/A [] Yes [] No []

Item	Feature	Subclause	Value/Comment	Status	Support
MGT9	Is the DRNI package supported?	Table 7–1, 7.4.1.1.1, 7.4.1.1.2, 7.4.1.1.3, 7.4.1.1.4, 7.4.1.1.5, 7.4.1.1.6, 7.4.1.1.7, 7.4.1.1.8, 7.4.1.1.9, 7.4.1.1.10, 7.4.1.1.11, 7.4.1.1.12, 7.4.1.1.13, 7.4.1.1.14, 7.4.1.1.15, 7.4.1.1.16, 7.4.1.1.20, 7.4.1.1.21, 7.4.1.1.22, 7.4.1.1.23, 7.4.2, 7.4.2.1.1, 7.4.2.1.2, 7.4.2.1.3, 7.4.2.1.4, 7.4.2.1.5, 7.4.2.1.6	Support of the Managed objects marked as members of the DRNI package in Table 7–1	MGT AND DRNI: M	N/A [] Yes []
MGT11	Is the aDrniEncapsulation-Method managed object supported?	7.4.1.1.17	This managed object is applicable only when Network / IPL sharing by time (9.3.2.1) or Network / IPL sharing by tag (9.3.2.2) or Network / IPL sharing by encapsulation (9.3.2.3) is supported	MGT AND DRNI9:M MGT AND DRNI10:M MGT AND DRNI11:M	N/A [] Yes []
MGT12	Is the aDrniIPLEncapMap managed object supported?	7.4.1.1.18	This managed object is applicable only when Network / IPL sharing by tag (9.3.2.2) or Network / IPL sharing by encapsulation (9.3.2.3) is supported	MGT AND DRNI10:M MGT AND DRNI11:M	N/A [] Yes []

Item	Feature	Subclause	Value/Comment	Status	Support
MGT13	Is the aDrniNetEncapMap managed object supported?	7.4.1.1.19	This managed object is applicable only when Network / IPL sharing by tag (9.3.2.2) is supported	MGT AND DRNI10:M	N/A [] Yes []
MGT14	Is the IPP statistics package supported?	Table 7-1, 7.4.3.1.1, 7.4.3.1.2, 7.4.3.1.3, 7.4.3.1.4	Support of the Managed objects marked as members of the IPP statistics package in Table 7-1	MGT AND DRNI: O	N/A [] Yes [] No []
MGT15	Is the IPP debug information package supported?	Table 7-1, 7.4.4.1.1, 7.4.4.1.2, 7.4.4.1.3, 7.4.4.1.4	Support of the Managed objects marked as members of the IPP debug information package in Table 7-1	MGT AND DRNI: O	N/A [] Yes [] No []

A.2.26 Per-Service Frame Distribution

Item	Feature	Subclause	Value/Comment	Status	Support
	If PSFD is not supported, mark N/A and ignore the remainder of this table				N/A []
PSFD1	Is Frame Distribution by VID supported?	8.2.2		PSFD: M	N/A [] Yes []
PSFD2	Is Frame Distribution by I-SID supported?	8.2.2, 8.2.3	Relevant only if Per-service frame distribution is implemented	PSFD: O	N/A [] Yes [] No []
PSFD3	Does the implementation support Frame Distribution by other methods not specified by this standard?			PSFD: O	N/A [] Yes [] No []

A.2.27 Conversation-sensitive frame collection and distribution

Item	Feature	Subclause	Value/Comment	Status	Support
	If CSCD is not supported, mark N/A and ignore the remainder of this table				N/A []

Item	Feature	Subclause	Value/Comment	Status	Support
CSCD1	Are the Conversation-sensitive collection and distribution state diagrams supported?	6.6.1	As shown in Figure 6-29 and as described	CSCD: M	Yes []
CSCD2	Does the function Determine-PortConversationID involve the application of local Service ID to Conversation ID mappings?	6.6.1.1.3, 7.3.1.1.38	Mapping ability required for classification by I-SID	CSCD:O CSCD AND PSFD2: M	Yes [] No []
CSCD3	Are the Verification state diagram (VER), the Report for Management Action state diagram (RMA), the Receive Long LACPDU state diagram (RXL) and the Update Mask state diagram (UM), the Actor CDS Churn Detection machine state diagram and the Partner CDS Churn Detection machine state diagram supported?	6.6.2	As shown in Figure 6-31, Figure 6-32, Figure 6-33, Figure 6-34, Figure 6-35, Figure 6-36, and as described	CSCD: O	Yes [] No []

A.2.28 Configuration capabilities and restrictions

Item	Feature	Subclause	Value/Comment	Status	Support
CCR1	Algorithm used to determine subset of Aggregation Ports that will be aggregated in Systems that have limited aggregation capability	6.7.1	As specified in items a) to e) of 6.7.1	M	Yes []
CCR2	Key value modification to generate optimum aggregation	6.7.2		O	Yes [] No []
CCR3	Key value modification when System has higher System Aggregation Priority			CCR2:M	N/A [] Yes []
CCR4	Key value modification when System has lower System Aggregation Priority			CCR2:X	N/A [] No []

A.2.29 Link Aggregation on shared-medium links

Item	Feature	Subclause	Value/Comment	Status	Support
LSM1	Shared-medium links— Configuration	6.7.3	Configured as Individual links	M	Yes []
LSM2	Shared-medium links— Operation of LACP	6.7.3	LACP is disabled	M	Yes []

A.2.30 Distributed Resilient Network Interconnect

Item	Feature	Subclause	Value/Comment	Status	Support
	If DRNI is not supported, mark N/A and ignore the remainder of this table				N/A []
DRNI1	Is the DR Function and the emulation of a Distributed Relay in cooperation with a single other Portal System, which also conforms to the provisions of this standard for DRNI, supported?	item b1) in 5.4, 9.2, 9.3	Support the state machines and procedures in 9.3 for a DR Function, constrained by the presence of up to a single IPL and up to two Portal Systems in the Portal, as specified	DRNI:M	Yes []
DRNI2	Is the DR Function and the emulation of a Distributed Relay, in cooperation with two other Portal Systems, which also conform to the provisions of this standard for DRNI, supported?	item a1) in 5.4.1, 9.2, 9.3	Support the state machines and procedures in 9.3 for a DR Function, for a Portal of 3 Portal Systems, as specified	DRNI:O	Yes [] No []
DRNI3	Is the Distributed Relay Control Protocol (DRCP) supported for Portal of two Portal Systems?	item b2) in 5.4, 9.4	Support the state machines and procedures in 9.4 as constrained by the presence of up to a single IPL and up to two Portal Systems in the Portal	DRNI:M	Yes []
DRNI4	Is the Distributed Relay Control Protocol (DRCP) supported for a Portal of three Portal Systems?	item a2) in 5.4.1, 9.4	Support the state machines and procedures in 9.4 for a Portal of 3 Portal Systems	DRNI2:M	Yes [] N/A []
DRNI5	Is a Portal Topology of three Portal Systems in a ring connected by three IPLs supported?	9.3.1, 9.4		DRNI4:O	Yes [] No [] N/A []
DRNI6	Are DRCPDUs supported?	item b3) in 5.4, 9.4.2	Transmit and receive DRCPDUs in the formats specified in 9.4.2 and 9.4.3	DRNI:M	Yes []
DRNI7	Can the IPL supported by using separate physical links for IPLs and for the network links?	item c) in 5.4, 9.3.2	Support separate physical links for IPLs and network links, as specified in item a)	DRNI:M	Yes []
DRNI8	Can the IPL be supported by a LAG and thus consisting of a number of physical links?	item b) in 5.4.1, 9.3.2	Support a separate Aggregation Port for the IPL, as specified in item b)	DRNI:O	Yes [] No []

Item	Feature	Subclause	Value/Comment	Status	Support
DRNI9	Can the IPL be supported by the Network / IPL sharing by time method?	item b) in 5.4.1, 9.3.2.1		DRNI:O	Yes [] No []
DRNI10	Can the IPL be supported by the Network / IPL sharing by tag method?	item b) in 5.4.1, 9.3.2.2		DRNI:O	Yes [] No []
DRNI11	Can the IPL be supported by the Network / IPL sharing by encapsulation method?	item b) in 5.4.1, 9.3.2.3		DRNI:O	Yes [] No []

A.2.31 DRCPDU structure

Item	Feature	Subclause	Value/Comment	Status	Support
	If DRNI6 is not supported, mark N/A and ignore the remainder of this table				N/A []
DRST1	DRCPDU addressing and protocol identifications	9.4.3	As described	DRNI6:M	Yes []
DRST2	DRCPDU structure	9.4.2	As shown in Figure 9-9 and as described	DRNI6:M	Yes []
DRST3	DRCPDU structure	9.4.3	All Reserved octets ignored on receipt and transmitted as zero	DRNI6:M	Yes []
DRST4	Are the Terminator TLV, the Portal Information TLV, the Portal Configuration the Information TLV, the DRCP State TLV, the Home Ports Information TLV, the Neighbor Ports Information TLV, the Home Gateway Vector TLV, the Neighbor Gateway Vector TLV and the Conversation Vector TLVs supported?	9.4.3, 9.4.3.2, 9.4.3.3	As shown in Figure 9-9, Figure 9-10, Figure 9-11, Figure 9-12, Figure 9-13, Figure 9-14 , Figure 9-15 , Figure 9-16 , Figure 9-17 , Figure 9-18 , and as described	DRNI6:M	Yes []
DRST5	Is the ONN bit in the Topology field within the Portal Configuration Information TLV supported?	item a3) in 5.4.1, 9.4.3	As shown in Figure 9-9 and as described	DRNI1:O DRNI2:M	Yes [] No [] N/A []
DRST6	Is the Other Gateway bit in the DRCP State TLV supported?	item a3) in 5.4.1, 9.4.3	As shown in Figure 9-9 and as described	DRNI1:O DRNI2:M	Yes [] No [] N/A []
DRST7	Is the Other Ports Information TLV and the Other Gateway Vector TLV supported?	item a3) in 5.4.1, 9.4.3	As shown in Figure 9-9 and as described	DRNI1:O DRNI2:M	Yes [] No [] N/A []

Item	Feature	Subclause	Value/Comment	Status	Support
DRST8	Is the Network/IPL Sharing Method TLV supported?	item b) in 5.4.1, 9.4.3.4.1	As shown in Figure 9-19 and as described	DRNI9:M DRNI10:M DRNI11:M	Yes [] N/A []
DRST9	Is the Network/IPL Sharing Encapsulation TLV supported?	item b) in 5.4.1, 9.4.3.4.2	As shown in Figure 9-20 and as described	DRNI10:M DRNI11:M	Yes [] N/A []
DRST10	Is the Organization-Specific TLV supported?	item c) in 5.4.1, 9.4.3.5	As shown in Figure 9-9 and as described	DRNI1:O DRNI2:M	Yes [] No [] N/A []

A.2.32 Bridge specific support

Item	Feature	Subclause	Value/Comment	Status	Support
BRG	Is the System implementing Link Aggregation is a IEEE Std 802.1Q Bridge? If not mark N/A and ignore the remainder of this table			O	Yes [] N/A []
BRG1	Is Per I-SID frame distribution supported on PIPs (Clause 6.10 of IEEE802.1Q-2011) or CBPs (Clause 6.11 of IEEE802.1Q-2011)	8.2.2, 8.2.3		BRG AND PSFD:M	Yes [] N/A []
BRG2	Is MAC Address Synchronization supported?	item b) in 5.4.1, Annex G	Only required when the IPL is supported by the Network / IPL sharing by time method and the Per Service frame distribution is not used	BRG AND DRNI9:M	Yes [] N/A []

Annex B

(informative)

Collection and distribution algorithms

B.1 Introduction

The specification of the Frame Collection and Frame Distribution functions was defined with the following considerations in mind:

- a) Frame duplication is not permitted.
- b) Frame ordering has to be preserved in aggregated links. Strictly, the Internal Sublayer Service specification (IEEE Std 802.1AC) states that order has to be preserved for frames with a given SA, DA, and priority; however, this is a tighter constraint than is absolutely necessary. There may be multiple, logically independent conversations in progress between a given SA-DA pair at a given priority; the real requirement is to maintain ordering within a conversation, though not necessarily between conversations.
- c) A single algorithm can be defined for the Frame Collection function that is independent of the distribution algorithm(s) employed by the Partner System.
- d) In the interests of simplicity and scalability, the Frame Collection function should not perform reassembly functions, reorder received frames, or modify received frames. Frame Distribution functions, therefore, do not make use of segmentation techniques, do not label or otherwise modify transmitted frames in any way, and have to operate in a manner that will inherently ensure proper ordering of received frames with the specified Frame Collector.
- e) The distribution and collection algorithms need to be capable of handling dynamic changes in aggregation membership.
- f) There are expected to be many different topologies and many different types of devices in which Link Aggregation will be employed. It is therefore unlikely that a single distribution algorithm will be applicable in all cases.

A simple Frame Collection function has been specified. The Frame Collector preserves the order of frames received on a given link, but does not preserve frame ordering among links. The Frame Distribution function maintains frame ordering by

- Transmitting frames of a given conversation on a single link at any time.
- Before changing the link on which frames of a given conversation are transmitted, ensuring that all previously transmitted frames of that conversation have been received to a point such that any subsequently transmitted frames received on a different links will be delivered to the Aggregator Client at a later time.

Given the wide variety of potential distribution algorithms, the normative text in Clause 6 specifies only the requirements that such algorithms have to meet, and not the details of the algorithms themselves. To clarify the intent, this informative annex gives examples of distribution algorithms, when they might be used, and the role of the Marker protocol (6.5) in their operation. The examples are not intended to be either exhaustive or prescriptive; implementers may make use of any distribution algorithms as long as the requirements of Clause 6 are met.

B.2 Port selection

A distribution algorithm selects the Aggregation Port used to transmit a given frame, such that the same Aggregation Port will be chosen for subsequent frames that form part of the same conversation. The algorithm may make use of information carried in the frame in order to make its decision, in combination with other information associated with the frame, such as its reception Aggregation Port in the case of a bridge.

The algorithm may assign one or more conversations to the same Aggregation Port; however, it has to not allocate some of the frames of a given conversation to one Aggregation Port and the remainder to different Aggregation Ports. The information used to assign conversations to Aggregation Ports could include the following:

- a) Source MAC address
- b) Destination MAC address
- c) The reception Aggregation Port
- d) The type of destination address (individual or group MAC address)
- e) Ethernet Length/Type value (i.e., protocol identification)
- f) Higher layer protocol information (e.g., addressing and protocol identification information from the LLC sublayer or above)
- g) Combinations of the above

One simple approach applies a hash function to the selected information to generate a Port Number. This produces a deterministic (i.e., history independent) Aggregation Port selection across a given number of Aggregation Ports in an aggregation. However, as it is difficult to select a hash function that will generate a uniform distribution of load across the set of Aggregation Ports for all traffic models, it might be appropriate to weight the Aggregation Port selection in favor of Aggregation Ports that are carrying lower traffic levels. In more sophisticated approaches, load balancing is dynamic; i.e., the Aggregation Port selected for a given set of conversations changes over time, independent of any changes that take place in the membership of the aggregation.

B.3 Dynamic reallocation of conversations to different Aggregation Ports

It may be necessary for a given conversation or set of conversations to be moved from one Aggregation Port to one or more others, as a result of

- a) An existing Aggregation Port being removed from the aggregation,
- b) A new Aggregation Port being added to the aggregation, or
- c) A decision on the part of the Frame Distributor to re-distribute the traffic across the set of Aggregation Ports.

Before moving conversation(s) to a new Aggregation Port, it is necessary to ensure that all frames already transmitted that are part of those conversations have been successfully received. The following procedure shows how the Marker protocol (6.5) can be used to ensure that no misordering of frames occurs:

- 1) Stop transmitting frames for the set of conversations affected. If the Aggregator Client requests transmission of further frames that are part of this set of conversations, these frames are discarded.
- 2) Start a timer, choosing the timeout period such that, if the timer expires, the destination System can be assumed either to have received or discarded all frames transmitted prior to starting the timer.
- 3) Use the Marker protocol to send a Marker PDU on the Aggregation Port previously used for this set of conversations.
- 4) Wait until either the corresponding Marker Response PDU is received or the timer expires.
- 5) Restart frame transmission for the set of conversations on the newly selected Aggregation Port.

1 The appropriate timeout value depends on the connected devices. For example, the recommended maximum
2 Bridge Transit Delay is 1 s; if the receiving device is a bridge, it may be expected to have forwarded or
3 discarded all frames received more than 1 s ago. The appropriate timeout value for other circumstances
4 could be smaller or larger than this by several orders of magnitude. For example, if the two Systems
5 concerned are high-performance end stations connected via Gigabit Ethernet links, then timeout periods
6 measured in milliseconds might be more appropriate. In order to allow an appropriate timeout value to be
7 determined, the Frame Collector parameter CollectorMaxDelay (see 6.2.3) defines the maximum delay that
8 the Frame Collector can introduce between receiving a frame from an Aggregation Port and either delivering
9 it to the Aggregator Client or discarding it. This value will be dependent upon the particular implementation
10 choices that have been made in a System. As far as the operation of the Frame Collector state machine is
11 concerned, CollectorMaxDelay is a constant; however, a management attribute, aAggCollectorMaxDelay
12 (7.3.1.1.32), is provided that allows interrogation and administrative control of its value. Hence, if a System
13 knows the value of CollectorMaxDelay that is in use by a Partner System, it can set the value of timeout
14 used when flushing a link to be equal to that value of CollectorMaxDelay, plus sufficient additional time to
15 allow for the propagation delay experienced by frames between the two Systems. A value of zero for the
16 CollectorMaxDelay parameter indicates that the delay imposed by the Frame Collector is less than the
17 resolution of the parameter (10 μ s). In this case, the delay that has to be considered is the physical
18 propagation delay of the channel. Allowing management manipulation of CollectorMaxDelay permits fine-
19 tuning of the value used in those cases where it may be difficult for the equipment to pre-configure a piece of
20 equipment with a realistic value for the physical propagation delay of the channel.

21
22 The Marker protocol provides an optimization that can result in faster reallocation of conversations than
23 would otherwise be possible—without the use of markers, the full timeout period would always have to be
24 used in order to be sure that no frames remained in transit between the local Frame Distributor and the
25 remote Frame Collector. The timeout described recovers from loss of Marker or Marker Response PDUs
26 that can occur.

27 28 29 **B.4 Topology considerations in the choice of distribution algorithm**

30
31 Figure B-1 gives some examples of different aggregated link scenarios. In some cases, it is possible to use
32 distribution algorithms that use MAC frame information to allocate conversations to links; in others, it is
33 necessary to make use of higher-layer information.

34
35 In example A, there is a many-to-many relationship between end stations communicating over the
36 aggregated link. It would be possible for each Bridge to allocate conversations to links simply on the basis of
37 source or destination MAC addresses.

38
39 In examples B and C, a number of end stations communicate with a single server via the aggregated link. In
40 these cases, the distribution algorithm employed in the server or in Bridge 2 can allocate traffic from the
41 server on the basis of destination MAC address; however, as one end of all conversations constitutes a single
42 server with a single MAC address, traffic from the end stations to the server would have to be allocated on
43 the basis of source MAC address. These examples illustrate the fact that different distribution algorithms can
44 be used in different devices, as appropriate to the circumstances. The collection algorithm is independent of
45 the distribution algorithm(s) that are employed.

46
47 In examples D and E, assuming that the servers are using a single MAC address for all of their traffic, the
48 only appropriate option is for the distribution algorithm used in the servers and Bridges to make use of
49 higher-layer information (e.g., Transport Layer socket identifiers) in order to allocate conversations to links.
50 Alternatively, in example E, if the servers were able to make use of multiple MAC addresses and allocate
51 conversations to them, then the Bridges could revert to MAC Address-based allocation.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

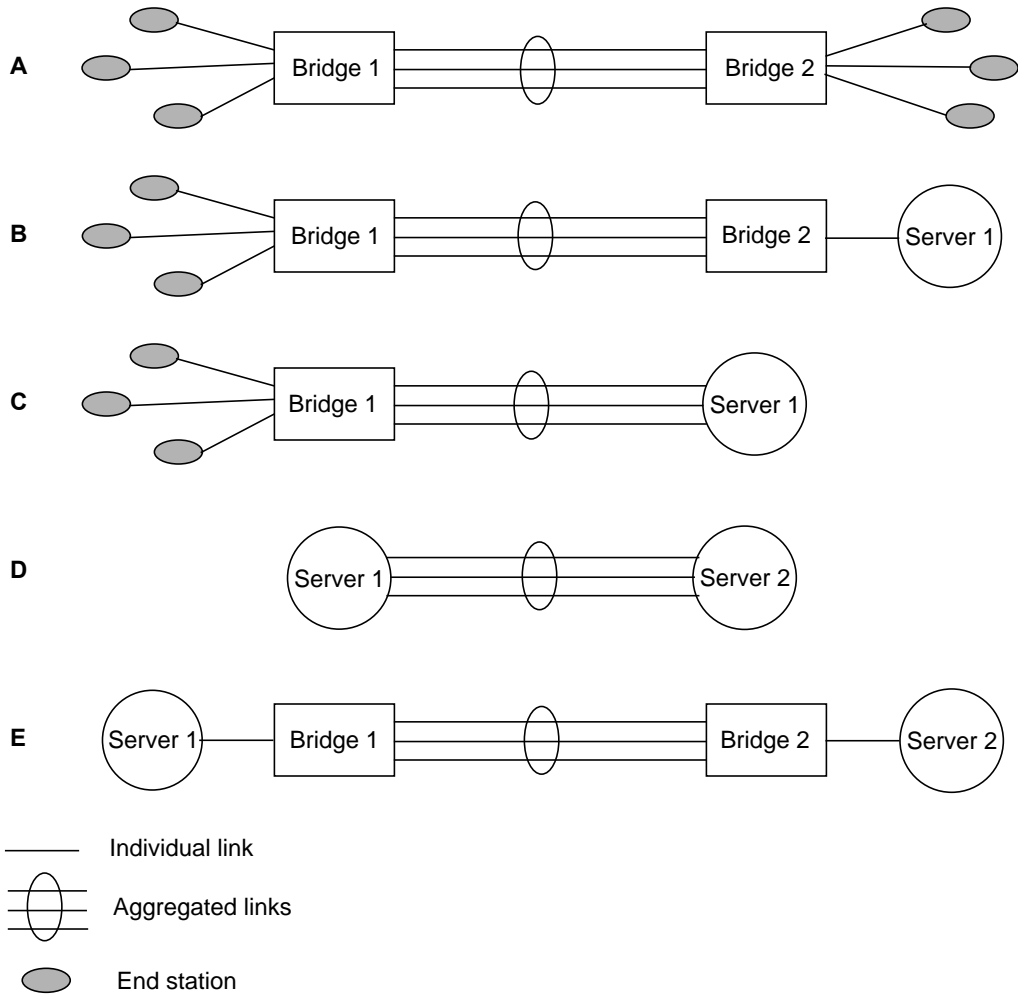


Figure B-1—Link aggregation topology examples

Annex C

(informative)

LACP standby link selection and dynamic Key management

C.1 Introduction

While any two Aggregation Ports on a given System that have been assigned the same administrative Key may be capable of aggregation, it is not necessarily the case that an arbitrary selection of such Aggregation Ports can be aggregated. (Keys may have been deliberately assigned to allow one link to be operated specifically as a hot standby for another.) A System may reasonably limit the number of Aggregation Ports attached to a single Aggregator, or the particular way more than two Aggregation Ports can be combined.

In cases where both communicating Systems have constraints on aggregation, it is necessary for them both to agree to some extent on the links to be selected for aggregation and on which not to use. Otherwise it might be possible for the two Systems to make different selections, possibly resulting in no communication at all.

When one or more links have to be selected as standby, it is possible that they could be used as part of a different Link Aggregation Group. For this to happen, one or another of the communicating Systems has to change the operational Key values used for the Aggregation Ports attached to those links.

If the operational Key values were to be changed independently by each System, the resulting set of aggregations could be unpredictable. It is possible that numerous aggregations, each containing a single link, may result. Worse, with no constraint on changes, the process of both Systems independently searching for the best combination of operational Key values may never end.

This annex describes protocol rules for standby link selection and dynamic Key management. It provides examples of a dynamic Key management algorithm applied to connections between Systems with various aggregation constraints.

C.2 Goals

The protocol rules presented

- a) Enable coordinated, predictable, and reproducible standby link selections.
- b) Permit predictable and reproducible partitioning of links into aggregations by dynamic Key management.

They do not require

- c) A LACP System to understand all the constraints on aggregations of multiple Aggregation Ports that might be imposed by other Systems.
- d) Correct configuration of parameters, i.e., they retain the plug and play attributes of LACP.

C.3 Standby link selection

Every link between Systems operating LACP is assigned a unique priority. This priority comprises (in priority order) the System Priority, System ID, Port Priority, and Port Number of the higher-priority System. In priority comparisons, numerically lower values have higher priority.

Aggregation Ports are considered for active use in an aggregation in link priority order, starting with the Aggregation Port attached to the highest priority link. Each Aggregation Port is selected for active use if preceding higher priority selections can also be maintained; otherwise, the Aggregation Port is selected as standby.

C.4 Dynamic Key management

Dynamic Key management changes the Key values used for links that either System has selected as a standby to allow use of more links. Whether this is desirable depends on their use. For example, if a single spanning tree is being used throughout the network, separating standby links into a separate aggregation serves little purpose. In contrast, if equal cost load sharing is being provided by routing, making additional bandwidth available in a separate Link Aggregation Group may be preferable to holding links in standby to provide link resilience.

The communicating System with the higher priority (as determined by System Priority and unique System ID) controls dynamic Key changes. Dynamic Key changes may only be made by this controlling System.

NOTE—The controlling System can observe the Port Priorities assigned by the Partner System, if it wishes to take these into account.

This rule prevents the confusion that could arise if both Systems change Keys simultaneously. In principle the controlling System might search all possible Key combinations for the best way to partition the links into groups. In practice the number of times that Keys may have to be changed to yield acceptable results is small.

After each Key change, the controlling System assesses which links are being held in standby by its Partner. Although there is no direct indication of this decision, standby links will be held OUT_OF_SYNC. After matched information is received from the protocol Partner, and before acting on this information, a “settling time” allows for the Partner’s aggregate wait delay, and for the selected links to be aggregated. Twice the Aggregate Wait Time (the expiry period for the wait_while_timer), i.e., 4 s, should be ample. If matched Partner information indicates that all the links that the Actor can make active have been brought IN_SYNC, it can proceed to change Keys on other links without further delay.

C.5 A dynamic Key management algorithm

The following algorithm is simple but effective.

After the “settling time” (see C.4) has elapsed, the controlling System scans its Aggregation Ports in the Link Aggregation Group (i.e., all those Aggregation Ports with a specific operational Key value that have the same Partner System Priority, System ID, and Key) in descending priority order.

For each Aggregation Port, it may wish to know

- a) Is the Aggregation Port (i.e., the Actor) *capable* of being aggregated with the Aggregation Ports already selected for aggregation with the current Key? Alternatively, is the Actor *not capable* of this aggregation?

- 1 b) Is the Aggregation Port's Partner IN_SYNC or is the Partner OUT_OF_SYNC?
2

3 As it inspects each Aggregation Port, it may
4

- 5 c) *Select* the Aggregation Port to be part of the aggregation with the current Key.
6 d) *Retain* the current Key for a further iteration of the algorithm, without selecting the Aggregation
7 Port to be part of the current aggregation.
8 e) *Change* the operational Key to a new value. Once a new value is chosen, all the Aggregation Ports in
9 the current Link Aggregation Group that have their Keys changed will be changed to this new value.
10

11 As the Aggregation Ports are scanned for the first time
12

- 13 1) The highest priority Aggregation Port is always selected.
14 If it is capable and IN_SYNC, move to step 2).
15 Otherwise, **change** the operational Key of all other Aggregation Ports (if any) in this Link
16 Aggregation Group, and apply this dynamic Key algorithm to those Aggregation Ports,
17 beginning with step 1), after the settling time.
18 2) Move to the next Aggregation Port.
19 If there is a next Aggregation Port, continue at step 3).
20 Otherwise, dynamic Key changes for Aggregation Ports with this operational Key are
21 complete.
22 Note that Aggregation Ports that were once in the same aggregation may have had their
23 operational Keys changed to (further) new values. If so, apply the dynamic Key management
24 algorithms to those Aggregation Ports, beginning with step 1), after the settling time.
25 3) If this Aggregation Port is capable and IN_SYNC:
26 **select** it, and repeat from step 2).
27 If this Aggregation Port is OUT_OF_SYNC:
28 **change** the operational Key, and repeat from step 2).
29 If this Aggregation Port is not capable but IN_SYNC:
30 **change** the operational Key, move to step 4).
31 4) Move to the next Aggregation Port.
32 If there is a next Aggregation Port, continue at step 5).
33 Otherwise If there are still Aggregation Ports in the current Link Aggregation Group (which
34 will have the current operational Key), wait for the settling time and apply the dynamic Key
35 management algorithm, beginning with the first such Aggregation Port, at step 3).
36 Otherwise, dynamic Key changes for Aggregation Ports with this operational Key are
37 complete.
38 5) **If** this Aggregation Port is capable:
39 retain the current Key and repeat from step 2).
40 Otherwise, **change** the operational Key and repeat from step 2).
41

42 This procedure is repeated until no OUT_OF_SYNC links remain, or a limit on the number of steps has been
43 reached.
44

45 If the Partner's System ID changes on any link at any time, the Actor's operational Key for that link should
46 revert to the administrative Key value, and the dynamic Key procedure should be rerun. This may involve
47 changing the operational Key values for all the links that were assigned Key values subsequent to the change
48 in Key for the link with the new Partner.
49

51 C.6 Example 1 52

53 Two Systems, A and B, are connected by four parallel links. Each System can support a maximum of two
54 links in an aggregation. They are connected as shown in Figure C-1. System A is the higher priority System.




Figure C-1—Example 1

The administrative Key for all of System A and System B's Aggregation Ports is 1. Neither System knows before the configuration is chosen that all its Aggregation Ports would attach to links of the same Partner System. Equally, if the links were attached to two different Systems, it is not known which pair of links (e.g., 1 and 2, or 1 and 4) would be attached to the same Partner. So choosing the administrative Keys values to be identical for four Aggregation Ports, even though only two could be actively aggregated, is very reasonable.

If there was no rule for selecting standby links, System A and System B might have both selected their own Aggregation Ports 1 and 2 as the active links, and there would be no communication. With the rule, the links A1-B4 and A2-B3 will become active, while A3-B2 and A4-B1 will be standby.

Since System A is the higher-priority System, System B's operational Key values will remain 1 while System A may dynamically change Keys, though it may choose to retain the standby links. Following the Key management algorithm suggested, System A would be able to change the Keys for A3 and A4 in a little over 2 seconds (depending on how fast System B completes the process of attaching its Aggregation Ports to the selected Aggregator) after the connections were first made, and both aggregations could be operating within 5 seconds.

If System A's aggregations were to be constrained to a maximum of three links, rather than two, while System B's are still constrained to two, the suggested algorithm would delay for 4 s before changing Keys. Both aggregations could be operating within 7 s.

C.7 Example 2

A System has the odd design constraint that each of its four Aggregation Ports may be aggregated with one other as follows:

- a) Aggregation Port 1 with Aggregation Port 2, or Aggregation Port 4.
- b) Aggregation Port 2 with Aggregation Port 3, or Aggregation Port 1.
- c) Aggregation Port 3 with Aggregation Port 4, or Aggregation Port 2.
- d) Aggregation Port 4 with Aggregation Port 1, or Aggregation Port 3.

This is equivalent to each Aggregation Port being able to aggregate with either neighbor, understanding the Aggregation Ports to be arranged in a circle.

Two such Systems are connected with four parallel links as shown in Figure C-2.




Figure C-2—Example 2a

13
14
15
16
17
18
19
20

Just as for Example 1, links A1-B4 and A2-B3 become active without changing the operational Key from its original administrative value. The Key for A3 and A4 is changed as soon as they become active, and a few seconds later A3-B2 and A4-B1 become active in a separate aggregation.

21
22

If the two Systems had been connected as shown in Figure C-3:




Figure C-3—Example 2b

33
34
35
36
37

the following would occur, assuming that System A operates the simple algorithm already described.

38
39
40
41
42
43
44

Initially System A advertises an operational Key equal to the administrative Key value of 1 on all Aggregation Ports. System B first selects B1 as active; since the link connects to A1 it has the highest priority. The next highest priority link is B3-A2, but System B cannot aggregate B3 with B1, so System B makes this Aggregation Port standby. System B can aggregate B4-A3, so the Aggregation Port is made active. Finally if B4 is aggregated with B1, B2 cannot be aggregated, so B2 is made standby.

45
46
47
48
49
50
51
52
53
54

System A, observing the resulting synchronization status from System B, assigns a Key value of 2 to Aggregation Ports 2 and 3, retaining the initial Key of 1 for Aggregation Ports 1 and 4. System B will remove B4 from the aggregation with B1, and substitute B2. B3 and B4 will be aggregated. In the final configuration A1-B1 and A4-B2 are aggregated, as are A2-B3 and A3-B4.

Annex D¹

(normative)

SMIv2 MIB definitions for Link Aggregation

D.1 Introduction

This annex defines a portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP based internets. In particular it defines objects for managing the operation of the Link Aggregation sublayer, based on the specification of Link Aggregation contained in this standard. This annex includes an MIB module that is SNMPv2 SMI compliant.

D.2 The SNMP Management Framework

For a detailed overview of the documents that describe the current Internet Standard Management Framework, please refer to section 7 of IETF RFC 3410 (2002).

Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. MIB objects are generally accessed through the Simple Network Management Protocol (SNMP). Objects in the MIB modules are defined using the mechanisms defined in the Structure of Management Information (SMI). This clause specifies a MIB module that is compliant to the SMIv2, which is described in IETF STD 58, comprising IETF RFC 2578 (1999), IETF RFC 2579 (1999), and IETF RFC 2580 (1999).

D.3 Security considerations

There are a number of management objects defined in this MIB module with a MAX-ACCESS clause of read-write and/or read-create. Such objects may be considered sensitive or vulnerable in some network environments. The support for SET operations in a non-secure environment without proper protection can have a negative effect on network operations.

This MIB module relates to systems that provide resilient transport at very high capacities based on link aggregation. Improper manipulation of the following tables and the objects in the tables may cause network instability and result in loss of service for a large number of end-users.

- dot3adAggTable
- dot3adAggXTable
- dot3adAggPortTable
- dot3adAggPortXTable
- dot3adAggPortSecondXTable

In particular, the object dot3adAggPortProtocolDA of the dot3adAggPortXTable sets the destination address for LACP frames and could be manipulated by an attacker to force a misconfiguration error causing the LAG to fail.

1. *Copyright release for SMIv2 MIB*: Users of this standard may freely reproduce the SMIv2 MIB in this annex so it can be used for its intended purpose.

1 Manipulation of the following objects can cause frames to be directed to the wrong ports, which can cause
2 loss of connectivity due to excessive congestion losses, as well as spoiling some or all of the goals of
3 Clause 8, such as bidirectional congruity.

4
5 dot3adAggConversationAdminPortTable
6 dot3adAggAdminServiceConversationMapTable
7

8 For a system implementing Distributed Resilient Network Interconnects, improper manipulation of the
9 following tables and the objects in the tables can have similar negative effects.

10
11 dot3adDrniTable
12 dot3adDrniConvAdminGatewayTable
13 dot3adDrniIPLEncapMapTable
14 dot3adDrniNetEncapMapTable
15 dot3adIPPAtributeTable
16

17 Some of the readable objects in this MIB module (i.e., objects with a MAX-ACCESS other than not-
18 accessible) may be considered sensitive or vulnerable in some network environments. It is thus important to
19 control even GET and/or NOTIFY access to these objects and possibly to even encrypt the values of these
20 objects when sending them over the network via SNMP. These are the tables whose objects contain
21 information which may be used to gain sensitive information which may be used to damage the organization
22 providing the link aggregation services or the users of those services. Sensitive information could be used by
23 an attacker to determine which attacks might be useful to attempt against a given device or to detect whether
24 attacks are being blocked or filtered.

25
26 dot3adAggPortStatsTable
27 dot3adAggPortDebugTable
28 dot3adAggPortDebugXTable
29 dot3adIPPDebugTable
30

31 SNMP versions prior to SNMPv3 did not include adequate security. Even if the network itself is secure (for
32 example by using IPsec), there is no control as to who on the secure network is allowed to access and GET/
33 SET (read/change/create/delete) the objects in this MIB module.

34
35 Implementations SHOULD provide the security features described by the SNMPv3 framework (see IETF
36 RFC 3410), and implementations claiming compliance to the SNMPv3 standard MUST include full support
37 for authentication and privacy via the User-based Security Model (USM) IETF RFC 3414 with the AES
38 cipher algorithm IETF RFC 3826. Implementations MAY also provide support for the Transport Security
39 Model (TSM) IETF RFC 5591 and the View-based Access Control Model IETF RFC 3415 in combination
40 with a secure transport such as SSH IETF RFC 5592 or TLS/DTLS IETF RFC 6353.

41
42 Further, deployment of SNMP versions prior to SNMPv3 is NOT RECOMMENDED. Instead, it is
43 RECOMMENDED to deploy SNMPv3 and to enable cryptographic security. It is then a customer/operator
44 responsibility to ensure that the SNMP entity giving access to an instance of this MIB module is properly
45 configured to give access to the objects only to those principals (users) that have legitimate rights to indeed
46 GET or SET (change/create/delete) them.

47 48 49 **D.4 Structure of the MIB module**

50
51 A single MIB module is defined in this annex. Objects in the MIB module are arranged into subtrees. Each
52 subtree is organized as a set of related objects. The overall structure and assignment of objects to their
53 subtrees and their corresponding definitions in Clause 7 is shown in the following subclauses.

D.4.1 Relationship to the managed objects defined in Clause 7

This subclause contains cross-references to the objects defined in Clause 7. Table D–1 contains cross references for the MIB module objects defined in this annex. Table D–2 contains definitions of ifTable elements, as defined in IETF RFC 2863, for an Aggregator. These table elements are cross referenced to the corresponding definitions in Clause 7.

Table D–1—Managed object cross reference table

Definition in Clause 7	MIB object
7.3.1.1.1 aAggID	ifIndex value
7.3.1.1.4 aAggActorSystemID	dot3adAggActorSystemID
7.3.1.1.5 aAggActorSystemPriority	dot3adAggActorSystemPriority
7.3.1.1.6 aAggAggregateOrIndividual	dot3adAggAggregateOrIndividual
7.3.1.1.7 aAggActorAdminKey	dot3adAggActorAdminKey
7.3.1.1.8 aAggActorOperKey	dot3adAggActorOperKey
7.3.1.1.10 aAggPartnerSystemID	dot3adAggPartnerSystemID
7.3.1.1.11 aAggPartnerSystemPriority	dot3adAggPartnerSystemPriority
7.3.1.1.12 aAggPartnerOperKey	dot3adAggPartnerOperKey
7.3.1.1.30 aAggPortList	dot3adAggPortListTable
7.3.1.1.32 aAggCollectorMaxDelay	dot3adAggCollectorMaxDelay
7.3.1.1.33 aAggPortAlgorithm	dot3adAggPortAlgorithm
7.3.1.1.34 aAggPartnerAdminPortAlgorithm	dot3adAggPartnerAdminPortAlgorithm
7.3.1.1.35 aAggConversationAdminLink[]	dot3adAggConversationAdminLinkTable
7.3.1.1.36 aAggPartnerAdminPortConversationListDigest	dot3adAggPartnerAdminPortConversationListDigest
7.3.1.1.37 aAggAdminDiscardWrongConversation	dot3adAggAdminDiscardWrongConversation
7.3.1.1.38 aAggAdminServiceConversationMap[]	dot3adAggAdminServiceConversationMapTable
7.3.1.1.39 aAggPartnerAdminConvServiceMappingDigest	dot3adAggPartnerAdminConvServiceMappingDigest
7.3.2.1.1 aAggPortID	ifIndex value
7.3.2.1.2 aAggPortActorSystemPriority	dot3adAggPortActorSystemPriority
7.3.2.1.3 aAggPortActorSystemID	dot3adAggPortActorSystemID
7.3.2.1.4 aAggPortActorAdminKey	dot3adAggPortActorAdminKey
7.3.2.1.5 aAggPortActorOperKey	dot3adAggPortActorOperKey
7.3.2.1.6 aAggPortPartnerAdminSystemPriority	dot3adAggPortPartnerAdminSystemPriority
7.3.2.1.7 aAggPortPartnerOperSystemPriority	dot3adAggPortPartnerOperSystemPriority
7.3.2.1.8 aAggPortPartnerAdminSystemID	dot3adAggPortPartnerAdminSystemID
7.3.2.1.9 aAggPortPartnerOperSystemID	dot3adAggPortPartnerOperSystemID
7.3.2.1.10 aAggPortPartnerAdminKey	dot3adAggPortPartnerAdminKey

Table D–1—Managed object cross reference table (continued)

Definition in Clause 7	MIB object
7.3.2.1.11 aAggPortPartnerOperKey	dot3adAggPortPartnerOperKey
7.3.2.1.12 aAggPortSelectedAggID	dot3adAggPortSelectedAggID
7.3.2.1.13 aAggPortAttachedAggID	dot3adAggPortAttachedAggID
7.3.2.1.14 aAggPortActorPort	dot3adAggPortActorPort
7.3.2.1.15 aAggPortActorPortPriority	dot3adAggPortActorPortPriority
7.3.2.1.16 aAggPortPartnerAdminPort	dot3adAggPortPartnerAdminPort
7.3.2.1.17 aAggPortPartnerOperPort	dot3adAggPortPartnerOperPort
7.3.2.1.18 aAggPortPartnerAdminPortPriority	dot3adAggPortPartnerAdminPortPriority
7.3.2.1.19 aAggPortPartnerOperPortPriority	dot3adAggPortPartnerOperPortPriority
7.3.2.1.20 aAggPortActorAdminState	dot3adAggPortActorAdminState
7.3.2.1.21 aAggPortActorOperState	dot3adAggPortActorOperState
7.3.2.1.22 aAggPortPartnerAdminState	dot3adAggPortPartnerAdminState
7.3.2.1.23 aAggPortPartnerOperState	dot3adAggPortPartnerOperState
7.3.2.1.24 aAggPortAggregateOrIndividual	dot3adAggPortAggregateOrIndividual
7.3.2.1.25 aAggPortOperConversationPasses	dot3adAggPortOperConversationPasses
7.3.2.1.26 aAggPortOperConversationCollected	dot3adAggPortOperConversationCollected
7.3.2.1.27 aAggPortLinkNumberID	dot3adAggPortLinkNumberId
7.3.2.1.28 aAggPortPartnerAdminLinkNumberID	dot3adAggPortPartnerAdminLinkNumberId
7.3.2.1.29 aAggPortWTRTime	dot3adAggPortWTRTime
7.3.2.2.1 aAggPortProtocolDA	dot3adAggPortProtocolDA
7.3.3.1.1 aAggPortStatsID	ifIndex value
7.3.3.1.2 aAggPortStatsLACPDUsRx	dot3adAggPortStatsLACPDUsRx
7.3.3.1.3 aAggPortStatsMarkerPDUsRx	dot3adAggPortStatsMarkerPDUsRx
7.3.3.1.4 aAggPortStatsMarkerResponsePDUsRx	dot3adAggPortStatsMarkerResponsePDUsRx
7.3.3.1.5 aAggPortStatsUnknownRx	dot3adAggPortStatsUnknownRx
7.3.3.1.6 aAggPortStatsIllegalRx	dot3adAggPortStatsIllegalRx
7.3.3.1.7 aAggPortStatsLACPDUsTx	dot3adAggPortStatsLACPDUsTx
7.3.3.1.8 aAggPortStatsMarkerPDUsTx	dot3adAggPortStatsMarkerPDUsTx
7.3.3.1.9 aAggPortStatsMarkerResponsePDUsTx	dot3adAggPortStatsMarkerResponsePDUsTx
7.3.4.1.1 aAggPortDebugInformationID	ifIndex value (see IETF RFC 2863) of the port
7.3.4.1.2 aAggPortDebugRxState	dot3adAggPortDebugRxState
7.3.4.1.3 aAggPortDebugLastRxTime	dot3adAggPortDebugLastRxTime
7.3.4.1.4 aAggPortDebugMuxState	dot3adAggPortDebugMuxState

Table D–1—Managed object cross reference table (continued)

Definition in Clause 7	MIB object
7.3.4.1.5 aAggPortDebugMuxReason	dot3adAggPortDebugMuxReason
7.3.4.1.6 aAggPortDebugActorChurnState	dot3adAggPortDebugActorChurnState
7.3.4.1.7 aAggPortDebugPartnerChurnState	dot3adAggPortDebugPartnerChurnState
7.3.4.1.8 aAggPortDebugActorChurnCount	dot3adAggPortDebugActorChurnCount
7.3.4.1.9 aAggPortDebugPartnerChurnCount	dot3adAggPortDebugPartnerChurnCount
7.3.4.1.10 aAggPortDebugActorSyncTransitionCount	dot3adAggPortDebugActorSyncTransitionCount
7.3.4.1.11 aAggPortDebugPartnerSyncTransitionCount	dot3adAggPortDebugPartnerSyncTransitionCount
7.3.4.1.12 aAggPortDebugActorChangeCount	dot3adAggPortDebugActorChangeCount
7.3.4.1.13 aAggPortDebugPartnerChangeCount	dot3adAggPortDebugPartnerChangeCount
7.3.4.1.14 aAggPortDebugActorCDSChurnState	dot3adAggPortDebugActorCDSChurnState
7.3.4.1.15 aAggPortDebugPartnerCDSChurnState	dot3adAggPortDebugPartnerCDSChurnState
7.3.4.1.16 aAggPortDebugActorCDSChurnCount	dot3adAggPortDebugActorCDSChurnCount
7.3.4.1.17 aAggPortDebugPartnerCDSChurnCount	dot3adAggPortDebugPartnerCDSChurnCount
7.4.1.1.1 aDrmiID	ifIndex value
7.4.1.1.2 aDrmiDescription	dot3adDrmiDescription
7.4.1.1.3 aDrmiName	dot3adDrmiName
7.4.1.1.4 aDrmiPortalAddr	dot3adDrmiPortalAddr
7.4.1.1.5 aDrmiPortalPriority	dot3adDrmiPortalPriority
7.4.1.1.6 aDrmiThreePortalSystem	dot3adDrmiThreePortalSystem
7.4.1.1.7 aDrmiPortalSystemNumber	dot3adDrmiPortalSystemNumber
7.4.1.1.8 aDrmiIntraPortalLinkList	dot3adDrmiIntraPortalLinkList
7.4.1.1.9 aDrmiAggregator	dot3adDrmiAggregator
7.4.1.1.10 aDrmiConvAdminGateway[]	dot3adDrmiConvAdminGatewayTable
7.4.1.1.11 aDrmiNeighborAdminConvGatewayListDigest	dot3adDrmiNeighborAdminConvGatewayListDigest
7.4.1.1.12 aDrmiNeighborAdminConvPortListDigest	dot3adDrmiNeighborAdminConvPortListDigest
7.4.1.1.13 aDrmiGatewayAlgorithm	dot3adDrmiGatewayAlgorithm
7.4.1.1.14 aDrmiNeighborAdminGatewayAlgorithm	dot3adDrmiNeighborAdminGatewayAlgorithm
7.4.1.1.15 aDrmiNeighborAdminPortAlgorithm	dot3adDrmiNeighborAdminPortAlgorithm
7.4.1.1.16 aDrmiNeighborAdminDRCPState	dot3adDrmiNeighborAdminDRCPState
7.4.1.1.17 aDrmiEncapsulationMethod	dot3adDrmiEncapsulationMethod
7.4.1.1.18 aDrmiIPLEncapMap	dot3adDrmiIPLEncapMapTable
7.4.1.1.19 aDrmiNetEncapMap	dot3adDrmiNetEncapMapTable
7.4.1.1.20 aDrmiDRPortConversationPasses	dot3adDrmiDRPortConversationPasses

Table D–1—Managed object cross reference table (continued)

Definition in Clause 7	MIB object
7.4.1.1.21 aDrmiDRGatewayConversationPasses	dot3adDrmiDRGatewayConversationPasses
7.4.1.1.22 aDrmiPSI	dot3adDrmiPSI
7.4.1.1.23 aDrmiPortConversationControl	dot3adDrmiPortConversationControl
7.4.1.1.24 aDrmiIntraPortalPortProtocolDA	dot3adDrmiIntraPortalPortProtocolDA
7.4.2.1.2 aIPPPortConversationPasses	dot3adIPPPortConversationPasses
7.4.2.1.3 aIPPGatewayConversationDirection	dot3adIPPGatewayConversationDirection
7.4.2.1.4 aIPPAAdminState	dot3adDrmiIPPAAdminState
7.4.2.1.5 aIPPOperState	dot3adDrmiIPPOperState
7.4.2.1.6 aIPPTimeOfLastOperChange	dot3adDrmiIPPTimeOfLastOperChange
7.4.3.1.1 aIPPStatsID	ifIndex value
7.4.3.1.2 aIPPStatsDRCPDUsRx	dot3adDrmiIPPStatsDRCPDUsRx
7.4.3.1.3 aIPPStatsIllegalRx	dot3adDrmiIPPStatsIllegalRx
7.4.3.1.4 aIPPStatsDRCPDUsTx	dot3adDrmiIPPStatsDRCPDUsTx
7.4.4.1.1 aIPPDebugInformationID	ifIndex value
7.4.4.1.2 aIPPDebugDRCPRxState	dot3adDrmiIPPDebugDRCPRxState
7.4.4.1.3 aIPPDebugLastRxTime	dot3adDrmiIPPDebugLastRxTime
7.4.4.1.4 aIPPDebugDifferPortalReason	dot3adIPPDebugDifferPortalReason

D.4.2 MIB Subtrees

The correspondence between the objects in the MIB module subtrees and the objects of the managed object classes defined in Clause 7 is described in the following subclauses.

D.4.2.1 The dot3adAgg Subtree

The objects of the dot3adAgg subtree correspond to the objects of the Aggregator managed object class (7.3.1) with the exception of the Aggregator Notifications (7.3.1.2).

D.4.2.2 The dot3adAggPort Subtree

The objects of the dot3adAggPort subtree correspond to the objects of the Aggregation Port managed object class (7.3.2), the Aggregation Port Statistics managed object class (7.3.3) and the Aggregation Port Debug Information managed object class (7.3.4).

D.4.2.3 The dot3adAggNotifications Subtree

The objects of the dot3adAggPort subtree correspond to the Aggregator Notifications (7.3.1.2).

Table D-2—ifTable element definitions for an Aggregator

Object	Definition
ifIndex	A unique integer value is allocated to each Aggregator by the local System. Interpreted as defined in IETF RFC 2863.
ifDescr	Interpreted as defined in IETF RFC 2863 and as further refined in the definition of aAggDescription (7.3.1.1.2).
ifType	ieee8023adLag(161) ^a .
ifMTU	The largest MAC Client SDU that can be carried by this Aggregator—1500 octets.
ifSpeed	The data rate of the Aggregation as defined for aAggDataRate (7.3.1.1.16).
ifPhysAddress	The individual MAC Address of the Aggregator as defined for aAggMACAddress (7.3.1.1.9).
ifAdminStatus	The administrative state of the Aggregator as defined for aAggAdminState (7.3.1.1.13).
ifOperStatus	The operational state of the Aggregator as defined for aAggOperState (7.3.1.1.14).
ifLastChange	Interpreted as defined in IETF RFC 2863; see also the definition of aAggTimeOfLastOperChange (7.3.1.1.15).
ifInOctets	The total number of user data octets received by the aggregation, as defined for aAggOctetsRxOK (7.3.1.1.18).
ifInUcastPkts	The total number of unicast user data frames received by the aggregation. This value is calculated as the value of aAggFramesRxOK (7.3.1.1.20), less the values of aAggMulticastFramesRxOK (7.3.1.1.22) and aAggBroadcastFramesRxOK (7.3.1.1.24).
ifInNUcastPkts	Deprecated in IETF RFC 2863.
ifInDiscards	The number of frames discarded on reception, as defined for aAggFramesDiscardedOnRx (7.3.1.1.26).
ifInErrors	The number of frames with reception errors, as defined for aAggFramesWithRxErrors (7.3.1.1.28).
ifInUnknownProtos	The number of unknown protocol frames discarded on reception, as defined for aAggUnknownProtocolFrames (7.3.1.1.29).
ifOutOctets	The total number of user data octets transmitted by the aggregation, as defined for aAggOctetsTxOK (7.3.1.1.17).
ifOutUcastPkts	The total number of unicast user data frames transmitted by the aggregation. This value is calculated as the value of aAggFramesTxOK (7.3.1.1.19), less the values of aAggMulticastFramesTxOK (7.3.1.1.21) and aAggBroadcastFramesTxOK (7.3.1.1.23).
ifOutNUcastPkts	Deprecated in IETF RFC 2863.

Table D–2—ifTable element definitions for an Aggregator (continued)

Object	Definition
ifOutDiscards	The number of frames discarded on transmission, as defined for aAggFramesDiscardedOnTx (7.3.1.1.25).
ifOutErrors	The number of frames discarded due to transmission errors, as defined for aAggFramesWithTxErrors (7.3.1.1.27).
ifOutQLen	Deprecated in IETF RFC 2863. Set to zero if present.
ifSpecific	Deprecated in IETF RFC 2863. Set to { 0.0 } if present.
ifLinkUpDownTrapEnable	See the definition of aAggLinkUpDownNotificationEnable (7.3.1.1.31).
ifConnectorPresent	“FALSE.”
ifHighSpeed	Set to zero.
ifName	The locally assigned textual name of the Aggregator, as defined for aAggName (7.3.1.1.3). Interpreted as defined in IETF RFC 2863.
linkUp TRAP	See the definition of nAggLinkUpNotification (7.3.1.2.1).
linkDown TRAP	See the definition of nAggLinkDownNotification (7.3.1.2.2).

a Values of ifType are assigned by the Internet Assigned Numbers Authority (IANA). A directory of number assignments is maintained on their website, at URL: <http://www.iana.org/numbers.html>. The currently assigned ifType values can be found in the SMI Numbers (Network Management Parameters) section of that directory.

D.4.2.4 The dot3adDrni Subtree

The objects of the dot3adDrni subtree correspond to the objects of the Distributed Relay Managed Object Class (7.4.1).

D.4.2.5 The dot3adIPP Subtree

The objects of the dot3adIPP subtree correspond to the objects of the IPP Managed Objects Class (7.4.2), the IPP Statistics managed object class (7.4.3) and the IPP Debug Information managed object class (7.4.4).

D.5 Relationship to other MIBs

It is assumed that a System implementing this MIB will also implement (at least) the “system” group defined in MIB-II defined in IETF RFC 1213 and the “interfaces” group defined in IETF RFC 2863.

D.5.1 Relationship to the Interfaces MIB

IETF RFC 2863, the Interface MIB Evolution, requires that any MIB that is an adjunct of the Interface MIB, clarify specific areas within the Interface MIB. These areas were intentionally left vague in IETF RFC 2863 to avoid over constraining the MIB, thereby precluding management of certain media types.

Section 3.3 of IETF RFC 2863 enumerates several areas that a media-specific MIB has to clarify. Each of these areas is addressed in D.5.2 and D.5.3. The implementer is referred to IETF RFC 2863 in order to understand the general intent of these areas.

1 In IETF RFC 2863, the “interfaces” group is defined as being mandatory for all Systems and contains
2 information on an entity’s interfaces, where each interface is thought of as being attached to a *subnetwork*.
3 (Note that this term is not to be confused with *subnet*, which refers to an addressing partitioning scheme
4 used in the Internet suite of protocols.) The term *segment* is sometimes used to refer to such a subnetwork.
5

6 Implicit in this MIB is the notion of Aggregators and Aggregation Ports. Each of these Aggregators and
7 Aggregation Ports is associated with one interface of the “interfaces” group (one row in the ifTable) and
8 each Aggregation Port is associated with a different interface.
9

10 Each Aggregator and Aggregation Port is uniquely identified by an interface number (ifIndex). The ifIndex
11 value assigned to a given Aggregation Port is the same as the ifIndex value assigned to the MAC interface
12 with which that Aggregation Port is associated.
13

14 **D.5.2 Layering model**

15 This annex assumes the interpretation of the Interfaces Group to be in accordance with IETF RFC 2863,
16 which states that the ifTable contains information on the managed resource’s interfaces and that each
17 sublayer below the internetwork layer of a network interface is considered an interface.
18

19 This annex recommends that, within an entity, aggregations that are instantiated as an entry in
20 dot3adAggTable are also represented by an entry in ifTable.
21

22 Where an entity contains Link Aggregation entities that transmit and receive traffic to/from an aggregation,
23 these should be represented in the ifTable as interfaces of type ieee8023adLag(161).
24

25 **D.5.3 ifStackTable**

26 If the ifStackTable defined in IETF RFC2863 is implemented, then
27

- 28 a) The relationship shown in the table has the property that an Aggregation is a higher interface relative
29 to an Aggregation Port.
- 30 b) This relationship is read-only.

31 NOTE—The restriction stated here is intended to enforce a strict hierarchical relationship between Aggregations and
32 Aggregation Ports, and to prevent those relationships from being modified. The read-only restriction does not apply to
33 any other relationships that may be expressed in the ifStackTable.
34

35 **D.5.4 ifRcvAddressTable**

36 The ifRcvAddressTable contains all MAC Addresses, unicast, multicast, and broadcast, for which an
37 interface can receive frames and forward them up to a higher layer entity for local consumption. An
38 Aggregator has at least one such address.
39

40 **D.6 Definitions for Link Aggregation MIB**

41 In the following MIB definition,² should there be any discrepancy between the DESCRIPTION text and the
42 BEHAVIOUR DEFINED AS in the corresponding definition in Clause 7, the definition in Clause 7 shall
43 take precedence.
44

45 _____
46
47
48
49
50
51
52
53
54 2. MIB definitions are available at <http://www.ieee802.org/1/pages/MIBS.html>.

1 NOTE—There is no requirement for persistency of the objects in the implementations of this MIB module. No guidance
2 is provided for the discontinuity of any counters in the MIB module.

3 NOTE 2—Many managed objects do not have DEFVAL values. Link Aggregation operation is designed to be
4 functional with any value of these parameters.
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

```

1  IEEE8023-LAG-MIB DEFINITIONS ::= BEGIN
2
3  -----
4  -- IEEE 802.1AX MIB from 802.1AXbk
5  -----
6
7
8  IMPORTS
9      MODULE-IDENTITY, OBJECT-TYPE, Counter32, Counter64, Integer32,
10         TimeTicks, NOTIFICATION-TYPE
11         FROM SNMPv2-SMI
12         DisplayString, MacAddress, TEXTUAL-CONVENTION, TruthValue
13         FROM SNMPv2-TC
14         MODULE-COMPLIANCE, OBJECT-GROUP, NOTIFICATION-GROUP
15         FROM SNMPv2-CONF
16         InterfaceIndex
17         FROM IF-MIB
18         PortList
19         FROM Q-BRIDGE-MIB
20         SnmpAdminString
21         FROM SNMP-FRAMEWORK-MIB
22         ;
23
24
25  lagMIB MODULE-IDENTITY
26      LAST-UPDATED "201409090000Z"
27      ORGANIZATION "IEEE 802.1 Working Group"
28      CONTACT-INFO
29          " WG-URL: http://grouper.ieee.org/groups/802/1/index.html
30            WG-EMail: stds-802-1@ieee.org
31
32            Contact: IEEE 802.1 Working Group Chair
33              Postal: C/O IEEE 802.1 Working Group
34                IEEE Standards Association
35                  445 Hoes Lane
36                    P.O. Box 1331
37                      Piscataway
38                        NJ 08855-1331
39                          USA
40                            E-mail: STDS-802-1-L@LISTSERV.IEEE.ORG"
41      DESCRIPTION
42          "The Link Aggregation module for managing IEEE 802.1AX-REV."
43      REVISION "201409090000Z"
44      DESCRIPTION
45          "The Link Aggregation module for managing IEEE 802.1AX."
46      REVISION "201201160000Z"
47      DESCRIPTION
48          "Updated for IEEE 802.1AXbk"
49      REVISION "200706290000Z"
50      DESCRIPTION
51          "References updated 04 Jun 2007 for IEEE 802.1AX"
52      REVISION "200006270000Z"
53      DESCRIPTION
54          "Original publication IEEE 802.3ad"

```

```

1      ::= { iso(1) member-body(2) us(840) 802dot3(10006) snmpmibs(300) 43 }
2
3
4      -- -----
5      -- Textual Conventions
6      -- -----
7
8
9      LACPKey ::= TEXTUAL-CONVENTION
10     DISPLAY-HINT "d"
11     STATUS      current
12     DESCRIPTION
13         "The Actor or Partner Key value."
14     SYNTAX      Integer32 (0..65535)
15
16     LACPState ::= TEXTUAL-CONVENTION
17     STATUS      current
18     DESCRIPTION
19         "The Actor and Partner State values from the LACPDU."
20     REFERENCE
21         "7.3.2.1.20"
22     SYNTAX      BITS {
23         lacpActivity(0),
24         lacpTimeout(1),
25         aggregation(2),
26         synchronization(3),
27         collecting(4),
28         distributing(5),
29         defaulted(6),
30         expired(7)
31     }
32
33     DRCPState ::= TEXTUAL-CONVENTION
34     STATUS      current
35     DESCRIPTION
36         "Administrative values of DRCP state."
37     SYNTAX      BITS {
38         homeGateway(0),
39         neighborGateway(1),
40         otherGateway(2),
41         ippActivity(3),
42         timeout(4),
43         gatewaySync(5),
44         portSync(6),
45         expired(7)
46     }
47
48     ChurnState ::= TEXTUAL-CONVENTION
49     STATUS      current
50     DESCRIPTION
51         "The state of the Churn Detection machine."
52     SYNTAX      INTEGER {
53         noChurn(1),
54         churn(2),

```

```

1           churnMonitor(3) -- deprecated
2     }
3
4   AggState ::= TEXTUAL-CONVENTION
5     STATUS      current
6     DESCRIPTION
7       "The state of the object entry."
8     SYNTAX      INTEGER {
9                 up(1),
10                down(2)
11            }
12
13   DrniConvAdminGatewayList ::= TEXTUAL-CONVENTION
14     DISPLAY-HINT "1x,"
15     STATUS      current
16     DESCRIPTION
17       "The three elements of the octet string represent the
18       three portal system numbers in order of priority with
19       highest priority first."
20     SYNTAX      OCTET STRING (SIZE (3))
21
22   PortalLinkList ::= TEXTUAL-CONVENTION
23     DISPLAY-HINT "4d,"
24     STATUS      current
25     DESCRIPTION
26       "Each four octets of the octet string represent an
27       ifIndex for an Intra-Port Link. The first ifIndex is
28       to Portal System 1, the second ifIndex is to Portal
29       System 2 and the third ifIndex is to portal System 3.
30       The ifIndex of the current portal system is set to zero."
31     SYNTAX      OCTET STRING (SIZE (12))
32
33   ServiceIdList ::= TEXTUAL-CONVENTION
34     DISPLAY-HINT "4d,"
35     STATUS      current
36     DESCRIPTION
37       "A list which contains, in general, a set of Service IDs
38       (8.2.2). If the Service IDs are representing VIDs, only a
39       single VID is applicable, while in the case that Service IDs
40       are representing I-SIDs, more than one I-SIDs are possible.
41       Each four octets represent a Service ID which may either be
42       I-SID or VID. An empty set is represented as an octet string
43       of size zero."
44     SYNTAX      OCTET STRING
45
46   -----
47   -- subtrees in the LAG MIB
48   -----
49
50   lagMIBNotifications OBJECT IDENTIFIER ::= { lagMIB 0 }
51   lagMIBObjects OBJECT IDENTIFIER      ::= { lagMIB 1 }
52   dot3adAggConformance OBJECT IDENTIFIER ::= { lagMIB 2 }
53
54   dot3adAgg OBJECT IDENTIFIER          ::= { lagMIBObjects 1 }

```



```

1 dot3adAggPort OBJECT IDENTIFIER ::= { lagMIBObjects 2 }
2 dot3adDrni OBJECT IDENTIFIER ::= { lagMIBObjects 4 }
3 dot3adIPP OBJECT IDENTIFIER ::= { lagMIBObjects 5 }
4
5 -----
6 -- The Tables Last Changed Object
7 -----
8
9 dot3adTablesLastChanged OBJECT-TYPE
10 SYNTAX TimeTicks
11 MAX-ACCESS read-only
12 STATUS current
13 DESCRIPTION
14 "This object indicates the time of the
15 most recent change to the dot3adAggTable,
16 dot3adAggPortTable, dot3adDrniTable and
17 dot3adIPPAttributeTable."
18
19 ::= { lagMIBObjects 3 }
20
21 -----
22 -- The Aggregator Configuration Table
23 -----
24
25 dot3adAggTable OBJECT-TYPE
26 SYNTAX SEQUENCE OF Dot3adAggEntry
27 MAX-ACCESS not-accessible
28 STATUS current
29 DESCRIPTION
30 "A table that contains information about every
31 Aggregator that is associated with this System."
32 REFERENCE
33 "7.3.1"
34 ::= { dot3adAgg 1 }
35
36 dot3adAggEntry OBJECT-TYPE
37 SYNTAX Dot3adAggEntry
38 MAX-ACCESS not-accessible
39 STATUS current
40 DESCRIPTION
41 "A list of the Aggregator parameters. This is indexed
42 by the ifIndex of the Aggregator."
43 INDEX { dot3adAggIndex }
44 ::= { dot3adAggTable 1 }
45
46 Dot3adAggEntry ::=
47 SEQUENCE {
48 dot3adAggIndex
49 InterfaceIndex,
50 dot3adAggMACAddress
51 MacAddress,
52 dot3adAggActorSystemPriority
53 Integer32,
54 dot3adAggActorSystemID

```

```

1         MacAddress,
2         dot3adAggAggregateOrIndividual
3         TruthValue,
4         dot3adAggActorAdminKey
5         LacpKey,
6         dot3adAggActorOperKey
7         LacpKey,
8         dot3adAggPartnerSystemID
9         MacAddress,
10        dot3adAggPartnerSystemPriority
11        Integer32,
12        dot3adAggPartnerOperKey
13        LacpKey,
14        dot3adAggCollectorMaxDelay
15        Integer32
16    }
17
18    dot3adAggIndex OBJECT-TYPE
19        SYNTAX      InterfaceIndex
20        MAX-ACCESS  not-accessible
21        STATUS      current
22        DESCRIPTION
23            "The unique identifier allocated to this Aggregator by the
24            local System. This attribute identifies an Aggregator instance
25            among the subordinate managed objects of the containing object.
26            This value is read-only. NOTE-The aAggID is represented in the
27            SMIV2 MIB as an ifIndex-see D.4.1."
28        REFERENCE
29            "7.3.1.1.1"
30        ::= { dot3adAggEntry 1 }
31
32
33    dot3adAggMACAddress OBJECT-TYPE
34        SYNTAX      MacAddress
35        MAX-ACCESS  read-only
36        STATUS      current
37        DESCRIPTION
38            "A 6-octet read-only value carrying the individual
39            MAC address assigned to the Aggregator."
40        REFERENCE
41            "7.3.1.1.9"
42        ::= { dot3adAggEntry 2 }
43
44
45    dot3adAggActorSystemPriority OBJECT-TYPE
46        SYNTAX      Integer32 (0..65535)
47        MAX-ACCESS  read-write
48        STATUS      current
49        DESCRIPTION
50            "A 2-octet read-write value indicating the priority value
51            associated with the Actor's System ID."
52        REFERENCE
53            "7.3.1.1.5"
54        ::= { dot3adAggEntry 3 }

```

1
2
3 dot3adAggActorSystemID OBJECT-TYPE
4 SYNTAX MacAddress
5 MAX-ACCESS read-write
6 STATUS current
7 DESCRIPTION
8 "A 6-octet read-write MAC address value used as a unique
9 identifier for the System that contains this Aggregator.
10 NOTE-From the perspective of the Link Aggregation
11 mechanisms described in Clause 6, only a single
12 combination of Actor's System ID and System Priority are
13 considered, and no distinction is made between the
14 values of these parameters for an Aggregator and the
15 port(s) that are associated with it; i.e., the protocol
16 is described in terms of the operation of aggregation
17 within a single System. However, the managed objects
18 provided for the Aggregator and the port both allow
19 management of these parameters. The result of this is to
20 permit a single piece of equipment to be configured by
21 management to contain more than one System from the
22 point of view of the operation of Link Aggregation. This
23 may be of particular use in the configuration of
24 equipment that has limited aggregation capability (see 6.7)."
25 REFERENCE
26 "7.3.1.1.4"
27 ::= { dot3adAggEntry 4 }
28
29
30 dot3adAggAggregateOrIndividual OBJECT-TYPE
31 SYNTAX TruthValue
32 MAX-ACCESS read-only
33 STATUS current
34 DESCRIPTION
35 "A read-only Boolean value indicating whether the
36 Aggregator represents an Aggregate ('TRUE') or
37 an Individual link ('FALSE')."
38 REFERENCE
39 "7.3.1.1.6"
40 ::= { dot3adAggEntry 5 }
41
42
43 dot3adAggActorAdminKey OBJECT-TYPE
44 SYNTAX LACPKey
45 MAX-ACCESS read-write
46 STATUS current
47 DESCRIPTION
48 "The current administrative value of the Key for the
49 Aggregator. The administrative Key value may differ from the
50 operational Key value for the reasons discussed in 6.7.2. This
51 is a 16-bit read-write value. The meaning of particular Key
52 values is of local significance. For an Aggregator that is
53 associated with a Portal the aAggActorAdminKey has to be
54 different for each Portal System. Specifically the two most

1 significant bits are set to aDrniPortalSystemNumber
2 (7.4.1.1.7). The lower 14 bits may be any value, have
3 to be the same in each Portal System within the same Portal,
4 and have a default of zero."
5 REFERENCE
6 "7.3.1.1.7"
7 ::= { dot3adAggEntry 6 }
8
9
10 dot3adAggActorOperKey OBJECT-TYPE
11 SYNTAX LacpKey
12 MAX-ACCESS read-only
13 STATUS current
14 DESCRIPTION
15 "The current operational value of the Key for the Aggregator.
16 The administrative Key value may differ from the operational
17 Key value for the reasons discussed in 6.7.2. This is a 16-bit
18 read-only value. The meaning of particular Key values is of
19 local significance."
20 REFERENCE
21 "7.3.1.1.8"
22 ::= { dot3adAggEntry 7 }
23
24
25 dot3adAggPartnerSystemID OBJECT-TYPE
26 SYNTAX MacAddress
27 MAX-ACCESS read-only
28 STATUS current
29 DESCRIPTION
30 "A 6-octet read-only MAC address value consisting of the
31 unique identifier for the current protocol Partner of this
32 Aggregator. A value of zero indicates that there is no known
33 Partner. If the aggregation is manually configured, this
34 System ID value will be a value assigned by the local System."
35 REFERENCE
36 "7.3.1.1.10"
37 ::= { dot3adAggEntry 8 }
38
39
40 dot3adAggPartnerSystemPriority OBJECT-TYPE
41 SYNTAX Integer32 (0..65535)
42 MAX-ACCESS read-only
43 STATUS current
44 DESCRIPTION
45 "A 2-octet read-only value that indicates the priority
46 value associated with the Partner's System ID. If the
47 aggregation is manually configured, this System Priority value
48 will be a value assigned by the local System."
49 REFERENCE
50 "7.3.1.1.11"
51 ::= { dot3adAggEntry 9 }
52
53
54 dot3adAggPartnerOperKey OBJECT-TYPE

```

1      SYNTAX      LacpKey
2      MAX-ACCESS  read-only
3      STATUS      current
4      DESCRIPTION
5          "The current operational value of the Key for the
6          Aggregator's current protocol Partner. This is a 16-bit
7          read-only value. If the aggregation is manually configured,
8          this Key value will be a value assigned by the local System."
9      REFERENCE
10         "7.3.1.1.12"
11         ::= { dot3adAggEntry 10 }
12
13
14 dot3adAggCollectorMaxDelay OBJECT-TYPE
15     SYNTAX      Integer32 (0..65535)
16     MAX-ACCESS  read-write
17     STATUS      current
18     DESCRIPTION
19         "The value of this 16-bit read-write attribute defines
20         the maximum delay, in tens of microseconds, that
21         may be imposed by the Frame Collector between
22         receiving a frame from an Aggregator Parser, and
23         either delivering the frame to its Aggregator Client
24         or discarding the frame (see 6.2.3.1.1)."

```

```

1
2 Dot3adAggPortListEntry ::=
3     SEQUENCE {
4         dot3adAggPortListPorts
5         PortList
6     }
7
8 dot3adAggPortListPorts OBJECT-TYPE
9     SYNTAX      PortList
10    MAX-ACCESS  read-only
11    STATUS      current
12    DESCRIPTION
13        "The complete set of ports currently associated with
14        this Aggregator. Each bit set in this list represents
15        an Actor Port member of this Link Aggregation."
16    REFERENCE
17        "7.3.1.1.30"
18    ::= { dot3adAggPortListEntry 1 }
19
20 -----
21 -- The Aggregation Extension Table
22 -----
23
24 dot3adAggXTable OBJECT-TYPE
25     SYNTAX      SEQUENCE OF Dot3adAggXEntry
26     MAX-ACCESS  not-accessible
27     STATUS      current
28     DESCRIPTION
29         "A table that extends dot3adAggTable."
30     REFERENCE
31         "7.3.1.1"
32     ::= { dot3adAgg 3 }
33
34 dot3adAggXEntry OBJECT-TYPE
35     SYNTAX      Dot3adAggXEntry
36     MAX-ACCESS  not-accessible
37     STATUS      current
38     DESCRIPTION
39         "A list of extension parameters for the Aggregator
40         Configuration Table"
41     AUGMENTS { dot3adAggEntry }
42     ::= { dot3adAggXTable 1 }
43
44 Dot3adAggXEntry ::=
45     SEQUENCE {
46         dot3adAggDescription
47         DisplayString,
48         dot3adAggName
49         DisplayString,
50         dot3adAggAdminState
51         AggState,
52         dot3adAggOperState
53         AggState,
54         dot3adAggTimeOfLastOperChange

```

```

1         Integer32,
2     dot3adAggDataRate
3         Integer32,
4     dot3adAggOctetsTxOK
5         Counter64,
6     dot3adAggOctetsRxOK
7         Counter64,
8     dot3adAggFramesTxOK
9         Counter64,
10    dot3adAggFramesRxOK
11        Counter64,
12    dot3adAggMulticastFramesTxOK
13        Counter64,
14    dot3adAggMulticastFramesRxOK
15        Counter64,
16    dot3adAggBroadcastFramesTxOK
17        Counter64,
18    dot3adAggBroadcastFramesRxOK
19        Counter64,
20    dot3adAggFramesDiscardedOnTx
21        Counter64,
22    dot3adAggFramesDiscardedOnRx
23        Counter64,
24    dot3adAggFramesWithTxErrors
25        Counter64,
26    dot3adAggFramesWithRxErrors
27        Counter64,
28    dot3adAggUnknownProtocolFrames
29        Counter64,
30    dot3adAggLinkUpDownNotificationEnable
31        Integer32,
32    dot3adAggPortAlgorithm
33        OCTET STRING,
34    dot3adAggPartnerAdminPortAlgorithm
35        OCTET STRING,
36    dot3adAggPartnerAdminPortConversationListDigest
37        OCTET STRING,
38    dot3adAggAdminDiscardWrongConversation
39        TruthValue,
40    dot3adAggPartnerAdminConvServiceMappingDigest
41        OCTET STRING
42    }
43
44    dot3adAggDescription OBJECT-TYPE
45        SYNTAX      DisplayString
46        MAX-ACCESS  read-only
47        STATUS      current
48        DESCRIPTION
49            "A human-readable text string containing information about
50            the Aggregator. This string could include information about
51            the distribution algorithm in use on this Aggregator; for
52            example, 'Aggregator 1, Dist Alg=Dest MAC address.' This
53            string is read-only. The contents are vendor specific."
54    REFERENCE

```

```
1         "7.3.1.1.2"
2         ::= { dot3adAggXEntry 1 }
3
4     dot3adAggName OBJECT-TYPE
5         SYNTAX      DisplayString
6         MAX-ACCESS  read-write
7         STATUS      current
8         DESCRIPTION
9             "A human-readable text string containing a locally significant
10            name for the Aggregator. This string is read-write."
11         REFERENCE
12            "7.3.1.1.3"
13         ::= { dot3adAggXEntry 2 }
14
15     dot3adAggAdminState OBJECT-TYPE
16         SYNTAX      AggState
17         MAX-ACCESS  read-write
18         STATUS      current
19         DESCRIPTION
20            "This read-write value defines the administrative state of
21            the Aggregator. A value of 'up' indicates that the operational
22            state of the Aggregator (aAggOperState) is permitted to be
23            either up or down. A value of 'down' forces the operational
24            state of the Aggregator to be down. Changes to the
25            administrative state affect the operational state of the
26            Aggregator only, not the operational state of the Aggregation
27            Ports that are attached to the Aggregator. A GET operation
28            returns the current administrative state. A SET operation
29            changes the administrative state to a new value."
30         REFERENCE
31            "7.3.1.1.13"
32         ::= { dot3adAggXEntry 3 }
33
34     dot3adAggOperState OBJECT-TYPE
35         SYNTAX      AggState
36         MAX-ACCESS  read-only
37         STATUS      current
38         DESCRIPTION
39            "This read-only value defines the operational state of the
40            Aggregator. An operational state of 'up' indicates that the
41            Aggregator is available for use by the Aggregator Client;
42            a value of 'down' indicates that the Aggregator is not
43            available for use by the Aggregator Client."
44         REFERENCE
45            "7.3.1.1.14"
46         ::= { dot3adAggXEntry 4 }
47
48     dot3adAggTimeOfLastOperChange OBJECT-TYPE
49         SYNTAX      Integer32
50         MAX-ACCESS  read-only
51         STATUS      current
52         DESCRIPTION
53            "The time at which the interface entered its current
54            operational state, in terms of centiseconds since the
```


1 system was last reset. If the current state was entered
2 prior to the last reinitialization of the local network
3 management subsystem, then this object contains a value of
4 zero. The ifLastChange object in the Interfaces MIB defined
5 in IETF RFC 2863 is a suitable object for supplying a value
6 for aAggTimeOfLastOperChange. This value is read-only."
7 REFERENCE
8 "7.3.1.1.15"
9 ::= { dot3adAggXEntry 5 }
10
11 dot3adAggDataRate OBJECT-TYPE
12 SYNTAX Integer32
13 MAX-ACCESS read-only
14 STATUS current
15 DESCRIPTION
16 "The current data rate, in bits per second, of the aggregate
17 link. The value is calculated as the sum of the data rate of
18 each link in the aggregation. This attribute is read-only."
19 REFERENCE
20 "7.3.1.1.16"
21 ::= { dot3adAggXEntry 6 }
22
23 dot3adAggOctetsTxOK OBJECT-TYPE
24 SYNTAX Counter64
25 UNITS
26 "octets"
27 MAX-ACCESS read-only
28 STATUS current
29 DESCRIPTION
30 "A count of the data and padding octets transmitted by this
31 Aggregator on all Aggregation Ports that are (or have been)
32 members of the aggregation. The count does not include octets
33 transmitted by the Aggregator in frames that carry LACPDUs or
34 Marker PDUs (7.3.3.1.7, 7.3.3.1.8, 7.3.3.1.9). However, it
35 includes frames discarded by the Frame Distribution function of
36 the Aggregator (7.3.1.1.25). This value is read-only."
37 REFERENCE
38 "7.3.1.1.17"
39 ::= { dot3adAggXEntry 7 }
40
41 dot3adAggOctetsRxOK OBJECT-TYPE
42 SYNTAX Counter64
43 UNITS
44 "octets"
45 MAX-ACCESS read-only
46 STATUS current
47 DESCRIPTION
48 "A count of the data and padding octets received by this
49 Aggregator, from the Aggregation Ports that are (or have been)
50 members of the aggregation. The count does not include octets
51 received in frames that carry LACP or Marker PDUs (7.3.3.1.2,
52 7.3.3.1.3, 7.3.3.1.4), or frames discarded by the Frame
53 Collection function of the Aggregator (7.3.1.1.26). This value
54 is read-only."

```
1      REFERENCE
2          "7.3.1.1.18"
3      ::= { dot3adAggXEntry 8 }
4
5  dot3adAggFramesTxOK OBJECT-TYPE
6      SYNTAX      Counter64
7      UNITS
8          "frames"
9      MAX-ACCESS  read-only
10     STATUS      current
11     DESCRIPTION
12         "A count of the data frames transmitted by this Aggregator on
13         all Aggregation Ports that are (or have been) members of the
14         aggregation. The count does not include frames transmitted by
15         the Aggregator that carry LACP or Marker PDUs (7.3.3.1.7,
16         7.3.3.1.8, 7.3.3.1.9). However, it includes frames discarded
17         by the Frame Distribution function of the Aggregator
18         (7.3.1.1.25). This value is read-only."
19     REFERENCE
20         "7.3.1.1.19"
21     ::= { dot3adAggXEntry 9 }
22
23  dot3adAggFramesRxOK OBJECT-TYPE
24     SYNTAX      Counter64
25     UNITS
26         "frames"
27     MAX-ACCESS  read-only
28     STATUS      current
29     DESCRIPTION
30         "A count of the data frames received by this Aggregator, from
31         the Aggregation Ports that are (or have been) members of the
32         aggregation. The count does not include frames that carry LACP
33         or Marker PDUs (7.3.3.1.2, 7.3.3.1.3, 7.3.3.1.4), or frames
34         discarded by the Frame Collection function of the Aggregator
35         (7.3.1.1.26). This value is read-only."
36     REFERENCE
37         "7.3.1.1.20"
38     ::= { dot3adAggXEntry 10 }
39
40  dot3adAggMulticastFramesTxOK OBJECT-TYPE
41     SYNTAX      Counter64
42     UNITS
43         "frames"
44     MAX-ACCESS  read-only
45     STATUS      current
46     DESCRIPTION
47         "A count of the data frames transmitted by this Aggregator on
48         all Aggregation Ports that are (or have been) members of the
49         aggregation, to a group destination address other than the
50         broadcast address. The count does not include frames
51         transmitted by the Aggregator that carry LACP or Marker PDUs
52         (7.3.3.1.7, 7.3.3.1.8, 7.3.3.1.9). However, it includes frames
53         discarded by the Frame Distribution function of the Aggregator
54         (7.3.1.1.25). This value is read-only."
```

```
1      REFERENCE
2          "7.3.1.1.21"
3      ::= { dot3adAggXEntry 11 }
4
5  dot3adAggMulticastFramesRxOK OBJECT-TYPE
6      SYNTAX      Counter64
7      UNITS
8          "frames"
9      MAX-ACCESS  read-only
10     STATUS      current
11     DESCRIPTION
12         "A count of the data frames received by this Aggregator, from
13         the Aggregation Ports that are (or have been) members of the
14         aggregation, that were addressed to an active group address
15         other than the broadcast address. The count does not include
16         frames that carry LACP or Marker PDUs (7.3.3.1.2, 7.3.3.1.3,
17         7.3.3.1.4), or frames discarded by the Frame Collection
18         function of the Aggregator (7.3.1.1.26). This value is
19         read-only."
20     REFERENCE
21         "7.3.1.1.22"
22     ::= { dot3adAggXEntry 12 }
23
24  dot3adAggBroadcastFramesTxOK OBJECT-TYPE
25     SYNTAX      Counter64
26     UNITS
27         "frames"
28     MAX-ACCESS  read-only
29     STATUS      current
30     DESCRIPTION
31         "A count of the broadcast data frames transmitted by this
32         Aggregator on all Aggregation Ports that are (or have been)
33         members of the aggregation. The count does not include frames
34         transmitted by the Aggregator that carry LACP or Marker PDUs
35         (7.3.3.1.7, 7.3.3.1.8, 7.3.3.1.9). However, it includes frames
36         discarded by the Frame Distribution function of the Aggregator
37         (7.3.1.1.25). This value is read-only."
38     REFERENCE
39         "7.3.1.1.23"
40     ::= { dot3adAggXEntry 13 }
41
42  dot3adAggBroadcastFramesRxOK OBJECT-TYPE
43     SYNTAX      Counter64
44     UNITS
45         "frames"
46     MAX-ACCESS  read-only
47     STATUS      current
48     DESCRIPTION
49         "A count of the broadcast data frames received by this
50         Aggregator, from the Aggregation Ports that are (or have been)
51         members of the aggregation. The count does not include frames
52         that carry LACP or Marker PDUs (7.3.3.1.2, 7.3.3.1.3,
53         7.3.3.1.4), illegal or unknown protocol frames (7.3.3.1.5,
54         7.3.3.1.6), or frames discarded by the Frame Collection
```

```
1         function of the Aggregator (7.3.1.1.26). This value is
2         read-only."
3     REFERENCE
4         "7.3.1.1.24"
5     ::= { dot3adAggXEntry 14 }
6
7     dot3adAggFramesDiscardedOnTx OBJECT-TYPE
8     SYNTAX      Counter64
9     UNITS
10        "frames"
11     MAX-ACCESS  read-only
12     STATUS      current
13     DESCRIPTION
14        "A count of data frames requested to be transmitted by this
15        Aggregator that were discarded by the Frame Distribution
16        function of the Aggregator when conversations are re-allocated
17        to different Aggregation Ports, due to the requirement to
18        ensure that the conversations are flushed on the old
19        Aggregation Ports in order to maintain proper frame ordering
20        (43A.3), or discarded as a result of excessive collisions by
21        Aggregation Ports that are (or have been) members of the
22        aggregation. This value is read-only."
23     REFERENCE
24        "7.3.1.1.25"
25     ::= { dot3adAggXEntry 15 }
26
27     dot3adAggFramesDiscardedOnRx OBJECT-TYPE
28     SYNTAX      Counter64
29     UNITS
30        "frames"
31     MAX-ACCESS  read-only
32     STATUS      current
33     DESCRIPTION
34        "A count of data frames, received on all Aggregation Ports
35        that are (or have been) members of the aggregation, that
36        were discarded by the Frame Collection function of the
37        Aggregator as they were received on Aggregation Ports whose
38        Frame Collection function was disabled. This value is
39        read-only."
40     REFERENCE
41        "7.3.1.1.26"
42     ::= { dot3adAggXEntry 16 }
43
44     dot3adAggFramesWithTxErrors OBJECT-TYPE
45     SYNTAX      Counter64
46     UNITS
47        "frames"
48     MAX-ACCESS  read-only
49     STATUS      current
50     DESCRIPTION
51        "A count of data frames requested to be transmitted by this
52        Aggregator that experienced transmission errors on Aggregation
53        Ports that are (or have been) members of the aggregation. This
54        count does not include frames discarded due to excess
```

```
1         collisions. This value is read-only."
2     REFERENCE
3         "7.3.1.1.27"
4     ::= { dot3adAggXEntry 17 }
5
6     dot3adAggFramesWithRxErrors OBJECT-TYPE
7         SYNTAX      Counter64
8         UNITS
9             "frames"
10        MAX-ACCESS  read-only
11        STATUS      current
12        DESCRIPTION
13            "A count of data frames discarded on reception by all
14            Aggregation Ports that are (or have been) members of the
15            aggregation, or that were discarded by the Frame Collection
16            function of the Aggregator, or that were discarded by the
17            Aggregator due to the detection of an illegal Slow Protocols
18            PDU (7.3.3.1.6). This value is read-only."
19        REFERENCE
20            "7.3.1.1.28"
21        ::= { dot3adAggXEntry 18 }
22
23        dot3adAggUnknownProtocolFrames OBJECT-TYPE
24            SYNTAX      Counter64
25            UNITS
26                "frames"
27            MAX-ACCESS  read-only
28            STATUS      current
29            DESCRIPTION
30                "A count of data frames discarded on reception by all
31                Aggregation Ports that are (or have been) members of the
32                aggregation, due to the detection of an unknown Slow Protocols
33                PDU (7.3.3.1.5). This value is read-only."
34            REFERENCE
35                "7.3.1.1.29"
36            ::= { dot3adAggXEntry 19 }
37
38        dot3adAggLinkUpDownNotificationEnable OBJECT-TYPE
39            SYNTAX      INTEGER {
40                enabled(1),
41                disabled(2)
42            }
43            MAX-ACCESS  read-write
44            STATUS      current
45            DESCRIPTION
46                "When set to 'enabled', Link Up and Link Down notifications are
47                enabled for this Aggregator. When set to 'disabled', Link Up
48                and Link Down notifications are disabled for this Aggregator.
49                This value is read-write."
50            REFERENCE
51                "7.3.1.1.31"
52            ::= { dot3adAggXEntry 20 }
53
54        dot3adAggPortAlgorithm OBJECT-TYPE
```

```
1      SYNTAX      OCTET STRING (SIZE (4))
2      MAX-ACCESS  read-write
3      STATUS      current
4      DESCRIPTION
5          "This object identifies the algorithm used by the Aggregator to
6          assign frames to a Port Conversation ID. Table 6-4 provides the
7          IEEE 802.1 OUI (00-80-C2) Port Algorithm encodings. A SEQUENCE
8          OF OCTETS consisting of a three-octet OUI or CID and one
9          following octet."
10     REFERENCE
11         "7.3.1.1.33"
12     ::= { dot3adAggXEntry 21 }
13
14 dot3adAggPartnerAdminPortAlgorithm OBJECT-TYPE
15     SYNTAX      OCTET STRING (SIZE (4))
16     MAX-ACCESS  read-write
17     STATUS      current
18     DESCRIPTION
19         "This object identifies the value for the algorithm of the
20         Partner System, assigned by administrator or System policy
21         for use when the Partner's information is unknown. Table 6-4
22         provides the IEEE 802.1 OUI (00-80-C2) Port Algorithm
23         encodings. Its default value is set to NULL. A SEQUENCE OF
24         OCTETS consisting of a three-octet OUI or CID and one following
25         octet."
26     REFERENCE
27         "7.3.1.1.34"
28     DEFVAL     { 'H' }
29     ::= { dot3adAggXEntry 22 }
30
31 dot3adAggPartnerAdminPortConversationListDigest OBJECT-TYPE
32     SYNTAX      OCTET STRING (SIZE (16))
33     MAX-ACCESS  read-write
34     STATUS      current
35     DESCRIPTION
36         "The value for the digest of the prioritized Port Conversation
37         ID-to-Link Number ID assignments of the Partner System, assigned
38         by administrator or System policy for use when the Partner's
39         information is unknown. Its default value is set to NULL."
40     REFERENCE
41         "7.3.1.1.36"
42     DEFVAL     { 'H' }
43     ::= { dot3adAggXEntry 23 }
44
45 dot3adAggAdminDiscardWrongConversation OBJECT-TYPE
46     SYNTAX      TruthValue
47     MAX-ACCESS  read-write
48     STATUS      current
49     DESCRIPTION
50         "The administrative value that determines what the Aggregator
51         does with a frame that is received from an Aggregation Port
52         with a Port Conversation ID that is not included in the
53         Collection_Conversation_Mask. The value 'TRUE' indicates that
54         such frames are to be discarded, and the value 'FALSE' that
```

```

1         they are to be forwarded. This variable needs to be set to
2         'TRUE', if bidirectional congruity (8.2.1) is required. Its
3         value is set to 'TRUE' by default."
4     REFERENCE
5         "7.3.1.1.37"
6     ::= { dot3adAggXEntry 24 }
7
8     dot3adAggPartnerAdminConvServiceMappingDigest OBJECT-TYPE
9     SYNTAX      OCTET STRING (SIZE (16))
10    MAX-ACCESS  read-write
11    STATUS      current
12    DESCRIPTION
13        "The value for the digest of the Port Conversation ID-to-Service
14        ID assignments of the Partner System, assigned by administrator
15        or System policy for use when the Partner's information is
16    unknown.
17        Its default value is set to NULL."
18    REFERENCE
19        "7.3.1.1.39"
20    DEFVAL     { 'H' }
21    ::= { dot3adAggXEntry 25 }
22
23    -----
24    -- The Aggregation Conversation Admin Link Table
25    -----
26
27    dot3adAggConversationAdminLinkTable OBJECT-TYPE
28    SYNTAX      SEQUENCE OF Dot3adAggConversationAdminLinkEntry
29    MAX-ACCESS  not-accessible
30    STATUS      current
31    DESCRIPTION
32        "There are 4096 aAggConversationAdminPort[] variables,
33        aAggConversationAdminLink[0] through
34        aAggConversationAdminLink[4095], indexed by Port Conversation
35        ID. Each contains administrative values of the link selection
36        priority list for the referenced Port Conversation ID. This
37        selection priority list is a sequence of Link-Number IDs for
38        each Port Conversation ID, in the order of preference, highest
39        to lowest, for the corresponding link to carry that Port
40        Conversation ID. A 16 bit zero value is used to indicate
41        that no link is assigned to carry the associated Port
42        Conversation ID. NOTE - This mapping of Port Conversation
43        IDs to Link Number IDs is the fundamental administrative input.
44        An equivalent mapping of Port Conversation IDs to Port IDs
45        [Conversation_PortList[] (6.6.2.1)] is derived from this and
46        used internally. NOTE - When a network administrator issues
47        a command for selection rules, provided by
48        aAggConversationAdminLink[], and accompanied with a non-zero
49        value for aAggPortWTRTime (7.3.2.1.29) for all associated
50        Aggregation Ports, the ChangeActorOperDist is set as specified
51        in 6.6.2.2. A value of 100 for the aAggPortWTRTime indicates
52        a non-revertive mode of operation and the WTR_timer will be
53        kept to the value 100."
54    REFERENCE

```

```

1         "7.3.1.1.35"
2         ::= { dot3adAgg 4 }
3
4     dot3adAggConversationAdminLinkEntry OBJECT-TYPE
5         SYNTAX      Dot3adAggConversationAdminLinkEntry
6         MAX-ACCESS  not-accessible
7         STATUS      current
8         DESCRIPTION
9             "An entry contains administrative values of the link selection
10            priority list for the referenced Port Conversation ID. This
11            selection priority list is a sequence of Link-Number IDs for
12            each Port Conversation ID, in the order of preference, highest
13            to lowest, for the corresponding link to carry that Port
14            Conversation ID. A 16 bit zero value is used to indicate
15            that no link is assigned to carry the associated Port
16            Conversation ID."
17         REFERENCE
18             "7.3.1.1.35"
19         INDEX { dot3adAggConversationAdminLinkId, dot3adAggIndex}
20         ::= { dot3adAggConversationAdminLinkTable 1 }
21
22
23     Dot3adAggConversationAdminLinkEntry ::=
24         SEQUENCE {
25             dot3adAggConversationAdminLinkId
26                 Integer32,
27             dot3adAggConversationAdminLinkList
28                 OCTET STRING
29         }
30
31     dot3adAggConversationAdminLinkId OBJECT-TYPE
32         SYNTAX      Integer32 (0..4095)
33         MAX-ACCESS  not-accessible
34         STATUS      current
35         DESCRIPTION
36             "An identifier for Port Conversation."
37         ::= { dot3adAggConversationAdminLinkEntry 1 }
38
39
40     dot3adAggConversationAdminLinkList OBJECT-TYPE
41         SYNTAX      OCTET STRING
42         MAX-ACCESS  read-write
43         STATUS      current
44         DESCRIPTION
45             "Each two octets of the octet string represent the agreed
46            Link Number ID that is assigned to an Aggregation Port
47            (7.3.2.1.27). The list is in in the order of preference,
48            highest to lowest, for corresponding preferred link to
49            carry that Port Conversation ID."
50         REFERENCE
51             "7.3.1.1.35"
52         ::= { dot3adAggConversationAdminLinkEntry 2 }
53
54     -----

```



```

1  -- The Aggregation Admin Service Conversation Map Table
2  -- -----
3
4  dot3adAggAdminServiceConversationMapTable OBJECT-TYPE
5      SYNTAX      SEQUENCE OF Dot3adAggAdminServiceConversationMapEntry
6      MAX-ACCESS  not-accessible
7      STATUS      current
8      DESCRIPTION
9          "There are 4096 aAggAdminServiceConversationMap[] variables,
10         aAggAdminServiceConversationMap[0] through
11         aAggAdminServiceConversationMap[4095], indexed by Port
12         Conversation ID. Each contains, in general, a set of Service
13         IDs (8.2.2), unique within the array. If the Service IDs are
14         representing VIDs, only a single VID is used, while in the case
15         that Service IDs are representing I-SIDs, more than one I-SIDs
16         are possible. Service IDs not contained in the map are not
17         mapped to any Port Conversation ID and will be discarded."
18      REFERENCE
19          "7.3.1.1.38"
20      ::= { dot3adAgg 5 }
21
22  dot3adAggAdminServiceConversationMapEntry OBJECT-TYPE
23      SYNTAX Dot3adAggAdminServiceConversationMapEntry
24      MAX-ACCESS  not-accessible
25      STATUS      current
26      DESCRIPTION
27          "An entry contains, in general, a set of Service IDs (8.2.2),
28          unique within the array. If the Service IDs are representing
29          VIDs, only a single VID is applicable, while in the case that
30          Service IDs are representing I-SIDs, more than one I-SIDs are
31          possible."
32      REFERENCE
33          "7.3.1.1.38"
34      INDEX { dot3adAggAdminServiceConversationMapId, dot3adAggIndex }
35      ::= { dot3adAggAdminServiceConversationMapTable 1 }
36
37
38
39  Dot3adAggAdminServiceConversationMapEntry ::=
40      SEQUENCE {
41          dot3adAggAdminServiceConversationMapId
42              Integer32,
43          dot3adAggAdminServiceConversationServiceIDList
44              ServiceIDList
45      }
46
47  dot3adAggAdminServiceConversationMapId OBJECT-TYPE
48      SYNTAX      Integer32 (0..4095)
49      MAX-ACCESS  not-accessible
50      STATUS      current
51      DESCRIPTION
52          "The Port Conversation ID used to index Conversation Map
53          entries."
54      ::= { dot3adAggAdminServiceConversationMapEntry 1}

```

```

1
2 dot3adAggAdminServiceConversationServiceIDList OBJECT-TYPE
3     SYNTAX      ServiceIdList
4     MAX-ACCESS  read-write
5     STATUS      current
6     DESCRIPTION
7         "A list contains, in general, a set of Service IDs (8.2.2),
8         unique within the array."
9     ::= { dot3adAggAdminServiceConversationMapEntry 2 }
10
11
12 -----
13 -- The Aggregation Port Table
14 -----
15
16 dot3adAggPortTable OBJECT-TYPE
17     SYNTAX      SEQUENCE OF Dot3adAggPortEntry
18     MAX-ACCESS  not-accessible
19     STATUS      current
20     DESCRIPTION
21         "A table that contains Link Aggregation Control
22         configuration information about every
23         Aggregation Port associated with this device.
24         A row appears in this table for each physical port."
25     REFERENCE
26         "7.3.2"
27     ::= { dot3adAggPort 1 }
28
29
30 dot3adAggPortEntry OBJECT-TYPE
31     SYNTAX      Dot3adAggPortEntry
32     MAX-ACCESS  not-accessible
33     STATUS      current
34     DESCRIPTION
35         "A list of Link Aggregation Control configuration
36         parameters for each Aggregation Port on this device."
37     INDEX       { dot3adAggPortIndex }
38     ::= { dot3adAggPortTable 1 }
39
40
41 Dot3adAggPortEntry ::=
42     SEQUENCE {
43         dot3adAggPortIndex
44         InterfaceIndex,
45         dot3adAggPortActorSystemPriority
46         Integer32,
47         dot3adAggPortActorSystemID
48         MacAddress,
49         dot3adAggPortActorAdminKey
50         LACPKey,
51         dot3adAggPortActorOperKey
52         LACPKey,
53         dot3adAggPortPartnerAdminSystemPriority
54         Integer32,

```

```

1      dot3adAggPortPartnerOperSystemPriority
2          Integer32,
3      dot3adAggPortPartnerAdminSystemID
4          MacAddress,
5      dot3adAggPortPartnerOperSystemID
6          MacAddress,
7      dot3adAggPortPartnerAdminKey
8          LACPKey,
9      dot3adAggPortPartnerOperKey
10         LACPKey,
11     dot3adAggPortSelectedAggID
12         InterfaceIndex,
13     dot3adAggPortAttachedAggID
14         InterfaceIndex,
15     dot3adAggPortActorPort
16         Integer32,
17     dot3adAggPortActorPortPriority
18         Integer32,
19     dot3adAggPortPartnerAdminPort
20         Integer32,
21     dot3adAggPortPartnerOperPort
22         Integer32,
23     dot3adAggPortPartnerAdminPortPriority
24         Integer32,
25     dot3adAggPortPartnerOperPortPriority
26         Integer32,
27     dot3adAggPortActorAdminState
28         LACPState,
29     dot3adAggPortActorOperState
30         LACPState,
31     dot3adAggPortPartnerAdminState
32         LACPState,
33     dot3adAggPortPartnerOperState
34         LACPState,
35     dot3adAggPortAggregateOrIndividual
36         TruthValue
37 }
38
39
40 dot3adAggPortIndex OBJECT-TYPE
41     SYNTAX      InterfaceIndex
42     MAX-ACCESS  not-accessible
43     STATUS      current
44     DESCRIPTION
45         "The unique identifier allocated to this Aggregation Port by
46         the local System. This attribute identifies an Aggregation
47         Port instance among the subordinate managed objects of the
48         containing object. This value is read-only. NOTE-The aAggPortID
49         is represented in the SMIV2 MIB as an ifIndex-see D.4.1."
50     REFERENCE
51         "7.3.2.1.1"
52     ::= { dot3adAggPortEntry 1 }
53
54

```

```
1 dot3adAggPortActorSystemPriority OBJECT-TYPE
2     SYNTAX      Integer32 (0..255)
3     MAX-ACCESS  read-write
4     STATUS      current
5     DESCRIPTION
6         "A 2-octet read-write value used to define the priority
7         value associated with the Actor's System ID."
8     REFERENCE
9         "7.3.2.1.2"
10    ::= { dot3adAggPortEntry 2 }
11
12
13 dot3adAggPortActorSystemID OBJECT-TYPE
14     SYNTAX      MacAddress
15     MAX-ACCESS  read-only
16     STATUS      current
17     DESCRIPTION
18         "A 6-octet read-only MAC address value that defines the
19         value of the System ID for the System that contains this
20         Aggregation Port."
21     REFERENCE
22         "7.3.2.1.3"
23    ::= { dot3adAggPortEntry 3 }
24
25
26 dot3adAggPortActorAdminKey OBJECT-TYPE
27     SYNTAX      LacpKey
28     MAX-ACCESS  read-write
29     STATUS      current
30     DESCRIPTION
31         "The current administrative value of the Key for the
32         Aggregation Port. This is a 16-bit read-write value.
33         The meaning of particular Key values is of local significance."
34     REFERENCE
35         "7.3.2.1.4"
36    ::= { dot3adAggPortEntry 4 }
37
38
39
40 dot3adAggPortActorOperKey OBJECT-TYPE
41     SYNTAX      LacpKey
42     MAX-ACCESS  read-only
43     STATUS      current
44     DESCRIPTION
45         "The current operational value of the Key for the
46         Aggregation Port. This is a 16-bit read-only value.
47         The meaning of particular Key values is of local significance."
48     REFERENCE
49         "7.3.2.1.5"
50    ::= { dot3adAggPortEntry 5 }
51
52
53
54 dot3adAggPortPartnerAdminSystemPriority OBJECT-TYPE
```

```
1      SYNTAX      Integer32 (0..255)
2      MAX-ACCESS  read-write
3      STATUS      current
4      DESCRIPTION
5          "A 2-octet read-write value used to define the administrative
6          value of priority associated with the Partner's System ID. The
7          assigned value is used, along with the value of
8          aAggPortPartnerAdminSystemID, aAggPortPartnerAdminKey,
9          aAggPortPartnerAdminPort, and aAggPortPartnerAdminPortPriority,
10         in order to achieve manually configured aggregation."
11     REFERENCE
12         "7.3.2.1.6"
13     ::= { dot3adAggPortEntry 6 }
14
15
16
17 dot3adAggPortPartnerOperSystemPriority OBJECT-TYPE
18     SYNTAX      Integer32 (0..255)
19     MAX-ACCESS  read-only
20     STATUS      current
21     DESCRIPTION
22         "A 2-octet read-only value indicating the operational value
23         of priority associated with the Partner's System ID. The
24         value of this attribute may contain the manually configured
25         value carried in aAggPortPartnerAdminSystemPriority if there
26         is no protocol Partner."
27     REFERENCE
28         "7.3.2.1.7"
29     ::= { dot3adAggPortEntry 7 }
30
31
32
33 dot3adAggPortPartnerAdminSystemID OBJECT-TYPE
34     SYNTAX      MacAddress
35     MAX-ACCESS  read-write
36     STATUS      current
37     DESCRIPTION
38         "A 6-octet read-write MACAddress value representing
39         the administrative value of the Aggregation Port's protocol
40         Partner's System ID. The assigned value is used, along with
41         the value of aAggPortPartnerAdminSystemPriority,
42         aAggPortPartnerAdminKey, aAggPortPartnerAdminPort,
43         and aAggPortPartnerAdminPortPriority, in order to
44         achieve manually configured aggregation."
45     REFERENCE
46         "7.3.2.1.8"
47     ::= { dot3adAggPortEntry 8 }
48
49
50 dot3adAggPortPartnerOperSystemID OBJECT-TYPE
51     SYNTAX      MacAddress
52     MAX-ACCESS  read-only
53     STATUS      current
54     DESCRIPTION
```

```
1         "A 6-octet read-only MACAddress value representing
2         the current value of the Aggregation Port's protocol Partner's
3         System ID. A value of zero indicates that there is no known
4         protocol Partner. The value of this attribute may contain the
5         manually configured value carried in
6         aAggPortPartnerAdminSystemID if there is no protocol Partner."
7     REFERENCE
8         "7.3.2.1.9"
9     ::= { dot3adAggPortEntry 9 }
10
11
12 dot3adAggPortPartnerAdminKey OBJECT-TYPE
13     SYNTAX      LacpKey
14     MAX-ACCESS  read-write
15     STATUS      current
16     DESCRIPTION
17         "The current administrative value of the Key for the
18         protocol Partner. This is a 16-bit read-write value.
19         The assigned value is used, along with the value of
20         aAggPortPartnerAdminSystemPriority,
21         aAggPortPartnerAdminSystemID, aAggPortPartnerAdminPort,
22         and aAggPortPartnerAdminPortPriority, in order to achieve
23         manually configured aggregation."
24     REFERENCE
25         "7.3.2.1.10"
26     ::= { dot3adAggPortEntry 10 }
27
28
29 dot3adAggPortPartnerOperKey OBJECT-TYPE
30     SYNTAX      LacpKey
31     MAX-ACCESS  read-only
32     STATUS      current
33     DESCRIPTION
34         "The current operational value of the Key for the
35         protocol Partner. The value of this attribute may contain
36         the manually configured value carried in
37         aAggPortPartnerAdminKey if there is no protocol Partner.
38         This is a 16-bit read-only value."
39     REFERENCE
40         "7.3.2.1.11"
41     ::= { dot3adAggPortEntry 11 }
42
43
44 dot3adAggPortSelectedAggID OBJECT-TYPE
45     SYNTAX      InterfaceIndex
46     MAX-ACCESS  read-only
47     STATUS      current
48     DESCRIPTION
49         "The identifier value of the Aggregator that this Aggregation
50         Port has currently selected. Zero indicates that the
51         Aggregation Port has not selected an Aggregator, either because
52         it is in the process of detaching from an Aggregator or because
53         there is no suitable Aggregator available for it to select.
54         This value is read-only."
```

```
1      REFERENCE
2          "7.3.2.1.12"
3      ::= { dot3adAggPortEntry 12 }
4
5
6      dot3adAggPortAttachedAggID OBJECT-TYPE
7          SYNTAX      InterfaceIndex
8          MAX-ACCESS  read-only
9          STATUS      current
10         DESCRIPTION
11             "The identifier value of the Aggregator that this Aggregation
12             Port is currently attached to. Zero indicates that the
13             Aggregation Port is not currently attached to an Aggregator.
14             This value is read-only."
15         REFERENCE
16             "7.3.2.1.13"
17         ::= { dot3adAggPortEntry 13 }
18
19
20         dot3adAggPortActorPort OBJECT-TYPE
21             SYNTAX      Integer32 (0..65535)
22             MAX-ACCESS  read-only
23             STATUS      current
24             DESCRIPTION
25                 "The port number locally assigned to the Aggregation Port.
26                 The port number is communicated in LACPDUs as the
27                 Actor_Port. This value is read-only."
28             REFERENCE
29                 "7.3.2.1.14"
30             ::= { dot3adAggPortEntry 14 }
31
32
33         dot3adAggPortActorPortPriority OBJECT-TYPE
34             SYNTAX      Integer32 (0..255)
35             MAX-ACCESS  read-write
36             STATUS      current
37             DESCRIPTION
38                 "The priority value assigned to this Aggregation Port.
39                 This 16-bit value is read-write. NOTE-In the case of DRNI
40                 (Clause 9), the two least significant bits of the priority
41                 for each Aggregation Port in a Distributed Relay's Aggregator
42                 Port will be ignored because these bits are used to encode
43                 the Portal System Number [item e) in 9.3.3]."
44             REFERENCE
45                 "7.3.2.1.15"
46             ::= { dot3adAggPortEntry 15 }
47
48
49         dot3adAggPortPartnerAdminPort OBJECT-TYPE
50             SYNTAX      Integer32 (0..65535)
51             MAX-ACCESS  read-write
52             STATUS      current
53             DESCRIPTION
54                 "The current administrative value of the port number
```

1 for the protocol Partner. This is a 16-bit read-write value.
2 The assigned value is used, along with the value of
3 aAggPortPartnerAdminSystemPriority,
4 aAggPortPartnerAdminSystemID, aAggPortPartnerAdminKey,
5 and aAggPortPartnerAdminPortPriority,
6 in order to achieve manually configured aggregation."

REFERENCE

7 "7.3.2.1.16"
8 ::= { dot3adAggPortEntry 16 }

dot3adAggPortPartnerOperPort OBJECT-TYPE

12 SYNTAX Integer32 (0..65535)
13 MAX-ACCESS read-only
14 STATUS current

DESCRIPTION

17 "The operational port number assigned to this Aggregation
18 Port by the Aggregation Port's protocol Partner. The value
19 of this attribute may contain the manually configured value
20 carried in aAggPortPartnerAdminPort if there is no protocol
21 Partner. This 16-bit value is read-only."

REFERENCE

22 "7.3.2.1.17"
23 ::= { dot3adAggPortEntry 17 }

dot3adAggPortPartnerAdminPortPriority OBJECT-TYPE

27 SYNTAX Integer32 (0..255)
28 MAX-ACCESS read-write
29 STATUS current

DESCRIPTION

31 "The current administrative value of the port priority
32 for the protocol Partner. This is a 16-bit read-write value.
33 The assigned value is used, along with the value of
34 aAggPortPartnerAdminSystemPriority,
35 aAggPortPartnerAdminSystemID, aAggPortPartnerAdminKey, and
36 aAggPortPartnerAdminPort, in order to achieve manually
37 configured aggregation."

REFERENCE

38 "7.3.2.1.18"
39 ::= { dot3adAggPortEntry 18 }

dot3adAggPortPartnerOperPortPriority OBJECT-TYPE

44 SYNTAX Integer32 (0..255)
45 MAX-ACCESS read-only
46 STATUS current

DESCRIPTION

48 "The priority value assigned to this Aggregation Port
49 by the Partner. The value of this attribute may contain the
50 manually configured value carried in
51 aAggPortPartnerAdminPortPriority if there is no
52 protocol Partner. This 16-bit value is read-only."

REFERENCE

53 "7.3.2.1.19"
54 ::= { dot3adAggPortEntry 19 }


```
1         "7.3.2.1.19"
2         ::= { dot3adAggPortEntry 19 }
3
4
5 dot3adAggPortActorAdminState OBJECT-TYPE
6     SYNTAX      LacpState
7     MAX-ACCESS  read-write
8     STATUS      current
9     DESCRIPTION
10        "A string of 8 bits, corresponding to the administrative
11        values of Actor_State (5.4.2) as transmitted by the Actor
12        in LACPDUs. The first bit corresponds to bit 0 of Actor_State
13        (LACP_Activity), the second bit corresponds to bit 1
14        (LACP_Timeout), the third bit corresponds to bit 2
15        (Aggregation), the fourth bit corresponds to bit 3
16        (Synchronization), the fifth bit corresponds to bit 4
17        (Collecting), the sixth bit corresponds to bit 5
18        (Distributing), the seventh bit corresponds to bit 6
19        (Defaulted), and the eighth bit corresponds to bit 7
20        (Expired). These values allow administrative control over
21        the values of LACP_Activity, LACP_Timeout and Aggregation.
22        This attribute value is read-write."
23     REFERENCE
24         "7.3.2.1.20"
25     ::= { dot3adAggPortEntry 20 }
26
27
28 dot3adAggPortActorOperState OBJECT-TYPE
29     SYNTAX      LacpState
30     MAX-ACCESS  read-only
31     STATUS      current
32     DESCRIPTION
33        "A string of 8 bits, corresponding to the current
34        operational values of Actor_State as transmitted by the
35        Actor in LACPDUs. The bit allocations are as defined in
36        7.3.2.1.20. This attribute value is read-only."
37     REFERENCE
38         "7.3.2.1.21"
39     ::= { dot3adAggPortEntry 21 }
40
41
42 dot3adAggPortPartnerAdminState OBJECT-TYPE
43     SYNTAX      LacpState
44     MAX-ACCESS  read-write
45     STATUS      current
46     DESCRIPTION
47        "A string of 8 bits, corresponding to the current
48        administrative value of Actor_State for the protocol Partner.
49        The bit allocations are as defined in 7.3.2.1.20. This
50        attribute value is read-write. The assigned value is used in
51        order to achieve manually configured aggregation."
52     REFERENCE
53         "7.3.2.1.22"
54     ::= { dot3adAggPortEntry 22 }
```

```
1
2
3 dot3adAggPortPartnerOperState OBJECT-TYPE
4     SYNTAX          LacpState
5     MAX-ACCESS      read-only
6     STATUS          current
7     DESCRIPTION
8         "A string of 8 bits, corresponding to the current values of
9         Actor_State in the most recently received LACPDU transmitted
10        by the protocol Partner. The bit allocations are as defined in
11        7.3.2.1.20. In the absence of an active protocol Partner, this
12        value may reflect the manually configured value
13        aAggPortPartnerAdminState. This attribute value is read-only."
14     REFERENCE
15         "7.3.2.1.23"
16     ::= { dot3adAggPortEntry 23 }
17
18
19 dot3adAggPortAggregateOrIndividual OBJECT-TYPE
20     SYNTAX          TruthValue
21     MAX-ACCESS      read-only
22     STATUS          current
23     DESCRIPTION
24         "A read-only Boolean value indicating whether the
25         Aggregation Port is able to Aggregate ('TRUE') or is
26         only able to operate as an Individual link ('FALSE')."
27     REFERENCE
28         "7.3.2.1.24"
29     ::= { dot3adAggPortEntry 24 }
30
31
32 -----
33 -- LACP Statistics Table
34 -----
35
36
37 dot3adAggPortStatsTable OBJECT-TYPE
38     SYNTAX          SEQUENCE OF Dot3adAggPortStatsEntry
39     MAX-ACCESS      not-accessible
40     STATUS          current
41     DESCRIPTION
42         "A table that contains Link Aggregation information
43         about every port that is associated with this device.
44         A row appears in this table for each physical port."
45     REFERENCE
46         "7.3.3"
47     ::= { dot3adAggPort 2 }
48
49
50 dot3adAggPortStatsEntry OBJECT-TYPE
51     SYNTAX          Dot3adAggPortStatsEntry
52     MAX-ACCESS      not-accessible
53     STATUS          current
54     DESCRIPTION
```

1 "A list of Link Aggregation Control Protocol statistics
2 for each port on this device."

3 INDEX { dot3adAggPortIndex }
4 ::= { dot3adAggPortStatsTable 1 }

5
6

7 Dot3adAggPortStatsEntry ::=

8 SEQUENCE {
9 dot3adAggPortStatsLACPDUsRx
10 Counter32,
11 dot3adAggPortStatsMarkerPDUsRx
12 Counter32,
13 dot3adAggPortStatsMarkerResponsePDUsRx
14 Counter32,
15 dot3adAggPortStatsUnknownRx
16 Counter32,
17 dot3adAggPortStatsIllegalRx
18 Counter32,
19 dot3adAggPortStatsLACPDUsTx
20 Counter32,
21 dot3adAggPortStatsMarkerPDUsTx
22 Counter32,
23 dot3adAggPortStatsMarkerResponsePDUsTx
24 Counter32
25 }

26
27

28 dot3adAggPortStatsLACPDUsRx OBJECT-TYPE

29 SYNTAX Counter32
30 MAX-ACCESS read-only
31 STATUS current
32 DESCRIPTION
33 "The number of valid LACPDUs received on this
34 Aggregation Port. This value is read-only."
35 REFERENCE
36 "7.3.3.1.2"
37 ::= { dot3adAggPortStatsEntry 1 }

38
39

40 dot3adAggPortStatsMarkerPDUsRx OBJECT-TYPE

41 SYNTAX Counter32
42 MAX-ACCESS read-only
43 STATUS current
44 DESCRIPTION
45 "The number of valid Marker PDUs received on this
46 Aggregation Port. This value is read-only."
47 REFERENCE
48 "7.3.3.1.3"
49 ::= { dot3adAggPortStatsEntry 2 }

50
51

52 dot3adAggPortStatsMarkerResponsePDUsRx OBJECT-TYPE

53 SYNTAX Counter32
54 MAX-ACCESS read-only

```
1      STATUS      current
2      DESCRIPTION
3          "The number of valid Marker Response PDUs received on this
4          Aggregation Port. This value is read-only."
5      REFERENCE
6          "7.3.3.1.4"
7      ::= { dot3adAggPortStatsEntry 3 }
8
9
10     dot3adAggPortStatsUnknownRx OBJECT-TYPE
11     SYNTAX      Counter32
12     MAX-ACCESS  read-only
13     STATUS      current
14     DESCRIPTION
15         "The number of frames received that either:
16         - carry the Slow Protocols Ethernet Type value
17         (IEEE Std 802.3 Annex 57A.4), but contain an
18         unknown PDU, or:
19         - are addressed to the Slow Protocols group MAC
20         Address (IEEE Std 802.3 Annex 57A.4), but do
21         not carry the Slow Protocols Ethernet Type.
22         This value is read-only."
23     REFERENCE
24         "7.3.3.1.5"
25     ::= { dot3adAggPortStatsEntry 4 }
26
27
28     dot3adAggPortStatsIllegalRx OBJECT-TYPE
29     SYNTAX      Counter32
30     MAX-ACCESS  read-only
31     STATUS      current
32     DESCRIPTION
33         "The number of frames received that carry the Slow
34         Protocols Ethernet Type value (IEEE Std 802.3 Annex
35         57A.4), but contain a badly formed PDU or an illegal
36         value of Protocol Subtype (IEEE Std 802.3 Annex 57A.4).
37         This value is read-only."
38     REFERENCE
39         "7.3.3.1.6"
40     ::= { dot3adAggPortStatsEntry 5 }
41
42
43     dot3adAggPortStatsLACPDUsTx OBJECT-TYPE
44     SYNTAX      Counter32
45     MAX-ACCESS  read-only
46     STATUS      current
47     DESCRIPTION
48         "The number of LACPDUs transmitted on this
49         Aggregation Port. This value is read-only."
50     REFERENCE
51         "7.3.3.1.7"
52     ::= { dot3adAggPortStatsEntry 6 }
53
54
```

```

1 dot3adAggPortStatsMarkerPDUsTx OBJECT-TYPE
2     SYNTAX          Counter32
3     MAX-ACCESS     read-only
4     STATUS         current
5     DESCRIPTION
6         "The number of Marker PDUs transmitted on this
7         Aggregation Port. This value is read-only."
8     REFERENCE
9         "7.3.3.1.8"
10    ::= { dot3adAggPortStatsEntry 7 }
11
12
13 dot3adAggPortStatsMarkerResponsePDUsTx OBJECT-TYPE
14     SYNTAX          Counter32
15     MAX-ACCESS     read-only
16     STATUS         current
17     DESCRIPTION
18         "The number of Marker Response PDUs transmitted
19         on this Aggregation Port. This value is read-only."
20     REFERENCE
21         "7.3.3.1.9"
22    ::= { dot3adAggPortStatsEntry 8 }
23
24
25 -----
26 -- LACP Debug Table
27 -----
28 dot3adAggPortDebugTable OBJECT-TYPE
29     SYNTAX          SEQUENCE OF Dot3adAggPortDebugEntry
30     MAX-ACCESS     not-accessible
31     STATUS         current
32     DESCRIPTION
33         "A table that contains Link Aggregation debug
34         information about every port that is associated with
35         this device. A row appears in this table for each
36         physical port."
37     REFERENCE
38         "7.3.4"
39    ::= { dot3adAggPort 3 }
40
41
42 dot3adAggPortDebugEntry OBJECT-TYPE
43     SYNTAX          Dot3adAggPortDebugEntry
44     MAX-ACCESS     not-accessible
45     STATUS         current
46     DESCRIPTION
47         "A list of the debug parameters for a port."
48     INDEX { dot3adAggPortIndex }
49    ::= { dot3adAggPortDebugTable 1 }
50
51
52 Dot3adAggPortDebugEntry ::=
53     SEQUENCE {
54         dot3adAggPortDebugRxState

```

```

1         Integer32,
2     dot3adAggPortDebugLastRxTime
3         TimeTicks,
4     dot3adAggPortDebugMuxState
5         Integer32,
6     dot3adAggPortDebugMuxReason
7         DisplayString,
8     dot3adAggPortDebugActorChurnState
9         ChurnState,
10    dot3adAggPortDebugPartnerChurnState
11        ChurnState,
12    dot3adAggPortDebugActorChurnCount
13        Counter32,
14    dot3adAggPortDebugPartnerChurnCount
15        Counter32,
16    dot3adAggPortDebugActorSyncTransitionCount
17        Counter32,
18    dot3adAggPortDebugPartnerSyncTransitionCount
19        Counter32,
20    dot3adAggPortDebugActorChangeCount
21        Counter32,
22    dot3adAggPortDebugPartnerChangeCount
23        Counter32
24    }

```

dot3adAggPortDebugRxState OBJECT-TYPE

```

28    SYNTAX          INTEGER {
29                currentRx(1),
30                expired(2),
31                defaulted(3),
32                initialize(4),
33                lacpDisabled(5),
34                portDisabled(6)
35    }

```

36 MAX-ACCESS read-only

37 STATUS current

38 DESCRIPTION

```

39        "This attribute holds the value 'currentRx' if the Receive
40        state machine for the Aggregation Port is in the
41        CURRENT state, 'expired' if the Receive state machine
42        is in the EXPIRED state, 'defaulted' if the Receive state
43        machine is in the DEFAULTED state, 'initialize' if the
44        Receive state machine is in the INITIALIZE state,
45        'lacpDisabled' if the Receive state machine is in the
46        LACP_DISABLED state, or 'portDisabled' if the Receive
47        state machine is in the PORT_DISABLED state.
48        This value is read-only."

```

49 REFERENCE

```

50        "7.3.4.1.2"

```

```

51    ::= { dot3adAggPortDebugEntry 1 }

```

52

53

54 dot3adAggPortDebugLastRxTime OBJECT-TYPE

```

1      SYNTAX      TimeTicks
2      MAX-ACCESS  read-only
3      STATUS      current
4      DESCRIPTION
5          "The value of aTimeSinceSystemReset (See IEEE Std 802.3 Annex
6          F.2.1) when the last LACPDU was received by this Aggregation
7          Port. This value is read-only."
8      REFERENCE
9          "7.3.4.1.3"
10     ::= { dot3adAggPortDebugEntry 2 }
11
12
13     dot3adAggPortDebugMuxState OBJECT-TYPE
14     SYNTAX      INTEGER {
15                 detached(1),
16                 waiting(2),
17                 attached(3),
18                 collecting(4),
19                 distributing(5),
20                 collectingDistributing(6)
21             }
22     MAX-ACCESS  read-only
23     STATUS      current
24     DESCRIPTION
25         "This attribute holds the value 'detached' if the Mux
26         state machine (5.4.14) for the Aggregation Port is
27         in the DETACHED state, 'waiting' if the Mux state machine
28         is in the WAITING state, 'attached' if the Mux state
29         machine for the Aggregation Port is in the ATTACHED
30         state, 'collecting' if the Mux state machine for the
31         Aggregation Port is in the COLLECTING state, 'distributing'
32         if the Mux state machine for the Aggregation Port is
33         in the DISTRIBUTING state, and 'collectingDistributing'
34         if the Mux state machine for the Aggregation Port is in
35         the COLLECTING_DISTRIBUTING state.
36         This value is read-only."
37     REFERENCE
38         "7.3.4.1.4"
39     ::= { dot3adAggPortDebugEntry 3 }
40
41
42     dot3adAggPortDebugMuxReason OBJECT-TYPE
43     SYNTAX      DisplayString
44     MAX-ACCESS  read-only
45     STATUS      current
46     DESCRIPTION
47         "A human-readable text string indicating the reason
48         for the most recent change of Mux machine state.
49         This value is read-only."
50     REFERENCE
51         "7.3.4.1.5"
52     ::= { dot3adAggPortDebugEntry 4 }
53
54

```

```
1 dot3adAggPortDebugActorChurnState OBJECT-TYPE
2     SYNTAX          ChurnState
3     MAX-ACCESS      read-only
4     STATUS          current
5     DESCRIPTION
6         "The state of the Actor Churn Detection machine
7         (6.4.17) for the Aggregation Port. A value of 'noChurn'
8         indicates that the state machine is in either the
9         NO_ACTOR_CHURN or the ACTOR_CHURN_MONITOR
10        state, and 'churn' indicates that the state machine is in the
11        ACTOR_CHURN state. This value is read-only."
12    REFERENCE
13        "7.3.4.1.6"
14    ::= { dot3adAggPortDebugEntry 5 }
15
16
17 dot3adAggPortDebugPartnerChurnState OBJECT-TYPE
18     SYNTAX          ChurnState
19     MAX-ACCESS      read-only
20     STATUS          current
21     DESCRIPTION
22        "The state of the Partner Churn Detection machine
23        (6.4.17) for the Aggregation Port. A value of 'noChurn'
24        indicates that the state machine is in either the
25        NO_PARTNER_CHURN or the PARTNER_CHURN_MONITOR
26        state, and 'churn' indicates that the state machine is
27        in the PARTNER_CHURN state.
28        This value is read-only."
29    REFERENCE
30        "7.3.4.1.7"
31    ::= { dot3adAggPortDebugEntry 6 }
32
33
34 dot3adAggPortDebugActorChurnCount OBJECT-TYPE
35     SYNTAX          Counter32
36     MAX-ACCESS      read-only
37     STATUS          current
38     DESCRIPTION
39        "Count of the number of times the Actor Churn state
40        machine has entered the ACTOR_CHURN state.
41        This value is read-only."
42    REFERENCE
43        "7.3.4.1.8"
44    ::= { dot3adAggPortDebugEntry 7 }
45
46
47 dot3adAggPortDebugPartnerChurnCount OBJECT-TYPE
48     SYNTAX          Counter32
49     MAX-ACCESS      read-only
50     STATUS          current
51     DESCRIPTION
52        "Count of the number of times the Partner Churn
53        state machine has entered the PARTNER_CHURN state.
54        This value is read-only."
```



```
1      REFERENCE
2          "7.3.4.1.9"
3      ::= { dot3adAggPortDebugEntry 8 }
4
5
6      dot3adAggPortDebugActorSyncTransitionCount OBJECT-TYPE
7          SYNTAX      Counter32
8          MAX-ACCESS  read-only
9          STATUS      current
10         DESCRIPTION
11             "Count of the number of times the Actor's Mux state
12             machine (6.4.15) has entered the IN_SYNC state.
13             This value is read-only."
14         REFERENCE
15             "7.3.4.1.10"
16         ::= { dot3adAggPortDebugEntry 9 }
17
18
19         dot3adAggPortDebugPartnerSyncTransitionCount OBJECT-TYPE
20             SYNTAX      Counter32
21             MAX-ACCESS  read-only
22             STATUS      current
23             DESCRIPTION
24                 "Count of the number of times the Partner's Mux
25                 state machine (6.4.15) has entered the IN_SYNC state.
26                 This value is read-only."
27             REFERENCE
28                 "7.3.4.1.11"
29             ::= { dot3adAggPortDebugEntry 10 }
30
31
32         dot3adAggPortDebugActorChangeCount OBJECT-TYPE
33             SYNTAX      Counter32
34             MAX-ACCESS  read-only
35             STATUS      current
36             DESCRIPTION
37                 "Count of the number of times the Actor's perception of
38                 the LAG ID for this Aggregation Port has changed.
39                 This value is read-only."
40             REFERENCE
41                 "7.3.4.1.12"
42             ::= { dot3adAggPortDebugEntry 11 }
43
44
45         dot3adAggPortDebugPartnerChangeCount OBJECT-TYPE
46             SYNTAX      Counter32
47             MAX-ACCESS  read-only
48             STATUS      current
49             DESCRIPTION
50                 "Count of the number of times the Partner's perception of
51                 the LAG ID (see 6.3.6.1) for this Aggregation Port has changed.
52                 This value is read-only."
53             REFERENCE
54                 "7.3.4.1.13"
```

```

1      ::= { dot3adAggPortDebugEntry 12 }
2
3
4
5      -----
6      -- Extension of the Aggregation Port Table
7      -----
8      dot3adAggPortXTable OBJECT-TYPE
9          SYNTAX      SEQUENCE OF Dot3adAggPortXEntry
10         MAX-ACCESS  not-accessible
11         STATUS      current
12         DESCRIPTION
13             "A table that extends dot3adAggPortTable."
14         REFERENCE
15             "7.3.2.2"
16         ::= { dot3adAggPort 4 }
17
18
19     dot3adAggPortXEntry OBJECT-TYPE
20         SYNTAX      Dot3adAggPortXEntry
21         MAX-ACCESS  not-accessible
22         STATUS      current
23         DESCRIPTION
24             "A list of extension parameters for Aggregation Port."
25         AUGMENTS { dot3adAggPortEntry }
26         ::= { dot3adAggPortXTable 1 }
27
28
29     Dot3adAggPortXEntry ::=
30         SEQUENCE {
31             dot3adAggPortProtocolDA
32             MacAddress
33         }
34
35     dot3adAggPortProtocolDA OBJECT-TYPE
36         SYNTAX      MacAddress
37         MAX-ACCESS  read-write
38         STATUS      current
39         DESCRIPTION
40             "A 6-octet read-write MACAddress value specifying the
41             destination address to be used when sending Link
42             Aggregation Control and Marker PDUs on this Aggregation
43             Port, corresponding to the value of Protocol_DA in
44             6.2.8.1.2, 6.2.10.1.3 and 6.5.4.2.1. The default value
45             shall be the the IEEE 802.3 Slow_Protocols_Multicast address."
46         REFERENCE
47             "7.3.2.2.1"
48         DEFVAL { '0180C2000002'H }
49         ::= { dot3adAggPortXEntry 1 }
50
51     -----
52     -- Second extension of the Aggregation Port Table
53     -----
54     dot3adAggPortSecondXTable OBJECT-TYPE

```

```

1      SYNTAX      SEQUENCE OF Dot3adAggPortSecondXEntry
2      MAX-ACCESS  not-accessible
3      STATUS      current
4      DESCRIPTION
5          "A table that extends dot3adAggPortTable."
6      REFERENCE
7          "7.3.2"
8      ::= { dot3adAggPort 5 }
9
10
11     dot3adAggPortSecondXEntry OBJECT-TYPE
12     SYNTAX      Dot3adAggPortSecondXEntry
13     MAX-ACCESS  not-accessible
14     STATUS      current
15     DESCRIPTION
16         "A list of extension parameters for Aggregation Port."
17     AUGMENTS { dot3adAggPortEntry }
18     ::= { dot3adAggPortSecondXTable 1 }
19
20     Dot3adAggPortSecondXEntry ::=
21     SEQUENCE {
22         dot3adAggPortOperConversationPasses
23             OCTET STRING,
24         dot3adAggPortOperConversationCollected
25             OCTET STRING,
26         dot3adAggPortLinkNumberId
27             Integer32,
28         dot3adAggPortPartnerAdminLinkNumberId
29             Integer32,
30         dot3adAggPortWTRTime
31             Integer32
32     }
33
34     dot3adAggPortOperConversationPasses OBJECT-TYPE
35     SYNTAX      OCTET STRING (SIZE (512))
36     MAX-ACCESS  read-only
37     STATUS      current
38     DESCRIPTION
39         "A read-only current operational vector of Boolean
40         values, with one value for each possible Port
41         Conversation ID. A 1 indicates that the Port
42         Conversation ID is distributed through this Aggregation
43         Port, and a 0 indicates that it cannot.
44         aAggPortOperConversationPasses is referencing the
45         current value of Port_Oper_Conversation_Mask (6.6.2.2)."

```

```
1         "A read-only current operational vector of Boolean values,
2         with one value for each possible Port Conversation ID. A 1
3         indicates that the Port Conversation ID is collected through
4         this Aggregation Port, and a 0 indicates that it cannot.
5         aAggPortOperConversationPasses is referencing the current
6         value of Collection_Conversation_Mask (6.6.1.1.2)."
```

REFERENCE

```
7
8         "7.3.2.1.26"
9         ::= { dot3adAggPortSecondXEntry 2 }
```

dot3adAggPortLinkNumberId OBJECT-TYPE

```
11
12     SYNTAX      Integer32 (0..65535)
13     MAX-ACCESS  read-write
14     STATUS      current
15     DESCRIPTION
```

"The Link Number ID value configured for this Aggregation Port by the System's administrator. When the Link Number ID value matches one of the non zero values in the selection prioritized lists in aAggConversationAdminLink[] (7.3.1.1.35), then this Aggregation Port must be configured to have an aAggPortActorAdminKey value that matches the aAggActorAdminKey of the Aggregator used by the LAG of the links specified in aAggConversationAdminLink[]. Its default value is set to aAggPortActorPort (7.3.2.1.14). NOTE - In the case of DRNI, the match of the aAggActorAdminKey to aAggPortActorAdminKey values excludes the first two bits identifying the individual Portal System in the Portal. If the network administrator fails to configure the proper values for the aAggPortActorAdminKey variables in all of the Aggregators Ports attached to a Portal, the Distributed Relay Control Protocol (DRCP, 9.4) and the variable Port_Oper_Conversation_Mask (6.6.2.2) prevent looping and/or duplicate delivery, if necessary, by discarding frames belonging to misconfigured Conversations."

REFERENCE

```
34
35     "7.3.2.1.27"
36     ::= { dot3adAggPortSecondXEntry 3 }
```

dot3adAggPortPartnerAdminLinkNumberId OBJECT-TYPE

```
38
39     SYNTAX      Integer32 (0..65535)
40     MAX-ACCESS  read-write
41     STATUS      current
42     DESCRIPTION
```

"The value for the Link Number ID of the Partner System for this Aggregation Port, assigned by administrator or System policy for use when the Partner's information is unknown. Its default value is set to 0."

REFERENCE

```
47
48     "7.3.2.1.28"
49     DEFVAL { 0 }
50     ::= { dot3adAggPortSecondXEntry 4 }
```

dot3adAggPortWTRTime OBJECT-TYPE

```
52
53     SYNTAX      Integer32 (0 | 5..12 | 100)
54     MAX-ACCESS  read-write
```

```

1      STATUS      current
2      DESCRIPTION
3          "The wait-to-restore (WTR) period accompanying selection
4          rules set by aAggConversationAdminLink[] in a command issued
5          by a network administrator. It may be configured in steps of
6          1 min between 5 min and 12 min, while two additional special
7          values are also used. The value 0 indicates revertive and is
8          the default value. The value 100 indicates non-revertive
9      mode
10         of operation and the WTR_timer will be kept to the value
11     100."
12     REFERENCE
13         "7.3.2.1.29"
14     DEFVAL { 0 }
15     ::= { dot3adAggPortSecondXEntry 5 }
16
17     -----
18     -- Extension of the LACP Debug Table
19     -----
20     dot3adAggPortDebugXTable OBJECT-TYPE
21         SYNTAX      SEQUENCE OF Dot3adAggPortDebugXEntry
22         MAX-ACCESS  not-accessible
23         STATUS      current
24         DESCRIPTION
25             "A table that extends dot3adAggPortDebugTable."
26         REFERENCE
27             "7.3.4"
28         ::= { dot3adAggPort 6 }
29
30
31     dot3adAggPortDebugXEntry OBJECT-TYPE
32         SYNTAX      Dot3adAggPortDebugXEntry
33         MAX-ACCESS  not-accessible
34         STATUS      current
35         DESCRIPTION
36             "A list of extension parameters for the LACP Port Debug table."
37         AUGMENTS { dot3adAggPortDebugEntry }
38         ::= { dot3adAggPortDebugXTable 1 }
39
40
41     Dot3adAggPortDebugXEntry ::=
42         SEQUENCE {
43             dot3adAggPortDebugActorCDSChurnState
44                 ChurnState,
45             dot3adAggPortDebugPartnerCDSChurnState
46                 ChurnState,
47             dot3adAggPortDebugActorCDSChurnCount
48                 Counter64,
49             dot3adAggPortDebugPartnerCDSChurnCount
50                 Counter64
51         }
52
53     dot3adAggPortDebugActorCDSChurnState OBJECT-TYPE
54         SYNTAX      ChurnState

```

```
1      MAX-ACCESS    read-only
2      STATUS        current
3      DESCRIPTION
4          "This managed object is applicable only when
5          Conversation-sensitive frame collection and distribution
6          as specified in 6.6 is supported. The state of the Actor
7          CDS Churn Detection machine (6.6.2.7) for the Aggregation
8          Port. A value of 'noChurn' indicates that the state machine
9          is in either the NO_ACTOR_CDS_CHURN or the
10         ACTOR_CHURN_CDS_MONITOR state, and 'churn' indicates that
11         the state machine is in the ACTOR_CDS_CHURN state. This value
12         is read-only."
13     REFERENCE
14         "7.3.4.1.14"
15     ::= { dot3adAggPortDebugXEntry 1 }
16
17 dot3adAggPortDebugPartnerCDSChurnState OBJECT-TYPE
18     SYNTAX          ChurnState
19     MAX-ACCESS      read-only
20     STATUS          current
21     DESCRIPTION
22         "This managed object is applicable only when
23         Conversation-sensitive frame collection and distribution
24         as specified in 6.6 is supported. The state of the Partner
25         CDS Churn Detection machine (6.6.2.7) for the Aggregation
26         Port. A value of 'noChurn' indicates that the state machine
27         is in either the NO_PARTNER_CDS_CHURN or the
28         PARTNER_CDS_CHURN_MONITOR state, and 'churn' indicates that
29         the state machine is in the PARTNER_CDSCHURN state. This value
30         is read-only."
31     REFERENCE
32         "7.3.4.1.15"
33     ::= { dot3adAggPortDebugXEntry 2 }
34
35 dot3adAggPortDebugActorCDSChurnCount OBJECT-TYPE
36     SYNTAX          Counter64
37     UNITS
38         "times entered ACTOR_CDS_CHURN"
39     MAX-ACCESS      read-only
40     STATUS          current
41     DESCRIPTION
42         "This managed object is applicable only when
43         Conversation-sensitive frame collection and distribution
44         as specified in 6.6 is supported. Count of the number of
45         times the Actor CDS Churn state machine has entered the
46         ACTOR_CDS_CHURN state. This value is read-only."
47     REFERENCE
48         "7.3.4.1.16"
49     ::= { dot3adAggPortDebugXEntry 3 }
50
51 dot3adAggPortDebugPartnerCDSChurnCount OBJECT-TYPE
52     SYNTAX          Counter64
53     UNITS
54         "times entered PARTNER_CDS_CHURN"
```

```

1      MAX-ACCESS    read-only
2      STATUS       current
3      DESCRIPTION
4          "This managed object is applicable only when
5          Conversation-sensitive frame collection and distribution
6          as specified in 6.6 is supported. Count of the number of
7          times the Partner CDS Churn state machine has entered the
8          PARTNER_CDS_CHURN state. This value is read-only."
9      REFERENCE
10         "7.3.4.1.7"
11         ::= { dot3adAggPortDebugXEntry 4 }
12
13
14     -----
15     -- The DRNI Configuration Table
16     -----
17
18     dot3adDrniTable OBJECT-TYPE
19         SYNTAX      SEQUENCE OF Dot3adDrniEntry
20         MAX-ACCESS  not-accessible
21         STATUS      current
22         DESCRIPTION
23             "A table that contains information about every
24             DRNI that is associated with this System."
25         REFERENCE
26             "7.4.1"
27         ::= { dot3adDrni 1 }
28
29     dot3adDrniEntry OBJECT-TYPE
30         SYNTAX      Dot3adDrniEntry
31         MAX-ACCESS  not-accessible
32         STATUS      current
33         DESCRIPTION
34             "A list of the DRNI parameters. This is indexed
35             by the DRNI Portal ID."
36         INDEX { dot3adDrniIndex }
37         ::= { dot3adDrniTable 1 }
38
39     Dot3adDrniEntry ::=
40         SEQUENCE {
41             dot3adDrniIndex
42             InterfaceIndex,
43             dot3adDrniDescription
44             SnmpAdminString,
45             dot3adDrniName
46             SnmpAdminString,
47             dot3adDrniPortalAddr
48             MacAddress,
49             dot3adDrniPortalPriority
50             Integer32,
51             dot3adDrniThreePortalSystem
52             TruthValue,
53             dot3adDrniPortalSystemNumber
54             Integer32,

```

```

1      dot3adDrniIntraPortalLinkList
2          PortalLinkList,
3      dot3adDrniAggregator
4          InterfaceIndex,
5      dot3adDrniNeighborAdminConvGatewayListDigest
6          OCTET STRING,
7      dot3adDrniNeighborAdminConvPortListDigest
8          OCTET STRING,
9      dot3adDrniGatewayAlgorithm
10         OCTET STRING,
11      dot3adDrniNeighborAdminGatewayAlgorithm
12         OCTET STRING,
13      dot3adDrniNeighborAdminPortAlgorithm
14         OCTET STRING,
15      dot3adDrniNeighborAdminDRCPState
16         DrcpState,
17      dot3adDrniEncapsulationMethod
18         OCTET STRING,
19      dot3adDrniDRPortConversationPasses
20         OCTET STRING,
21      dot3adDrniDRGatewayConversationPasses
22         OCTET STRING,
23      dot3adDrniPSI
24         TruthValue,
25      dot3adDrniPortConversationControl
26         TruthValue,
27      dot3adDrniIntraPortalPortProtocolDA
28         MacAddress
29  }
30
31  dot3adDrniIndex OBJECT-TYPE
32      SYNTAX      InterfaceIndex
33      MAX-ACCESS  not-accessible
34      STATUS      current
35      DESCRIPTION
36          "The unique identifier allocated to this Distributed Relay by
37          the local System. This attribute identifies a Distributed Relay
38          instance among the subordinate managed objects of the
39          containing object. This value is read-only. NOTE - The aDrniID
40          is represented in the SMIV2 MIB as an ifIndex-see D.5."
41      REFERENCE
42          "7.4.1.1.1"
43      ::= { dot3adDrniEntry 1 }
44
45  dot3adDrniDescription OBJECT-TYPE
46      SYNTAX      SnmpAdminString
47      MAX-ACCESS  read-only
48      STATUS      current
49      DESCRIPTION
50          "A human-readable text string containing information about the
51          Distribute Relay. This string is read-only. The contents are
52          vendor specific."
53      REFERENCE
54          "7.4.1.1.2"

```



```
1      ::= { dot3adDrniEntry 2 }
2
3  dot3adDrniName OBJECT-TYPE
4      SYNTAX      SnmpAdminString
5      MAX-ACCESS  read-write
6      STATUS      current
7      DESCRIPTION
8          "A human-readable text string containing a locally significant
9          name for the Distributed Relay. This string is read-write."
10     REFERENCE
11         "7.4.1.1.3"
12     ::= { dot3adDrniEntry 3 }
13
14  dot3adDrniPortalAddr OBJECT-TYPE
15     SYNTAX      MacAddress
16     MAX-ACCESS  read-write
17     STATUS      current
18     DESCRIPTION
19         "A read-write identifier of a particular Portal.
20         aDrniPortalAddr has to be unique among at least all of the
21         potential Portal Systems to which a given Portal System might
22         be attached via an IPL. Also used as the Actor's System ID
23         (6.3.2) for the emulated system."
24     REFERENCE
25         "7.4.1.1.4"
26     ::= { dot3adDrniEntry 4 }
27
28
29  dot3adDrniPortalPriority OBJECT-TYPE
30     SYNTAX      Integer32
31     MAX-ACCESS  read-write
32     STATUS      current
33     DESCRIPTION
34         "A 2-octet read-write value indicating the priority value
35         associated with the Portal's System ID. Also used as the
36         Actor's System Priority (6.3.2) for the emulated system."
37     REFERENCE
38         "7.4.1.1.5"
39     ::= { dot3adDrniEntry 5 }
40
41  dot3adDrniThreePortalSystem OBJECT-TYPE
42     SYNTAX      TruthValue
43     MAX-ACCESS  read-write
44     STATUS      current
45     DESCRIPTION
46         "A read-write Boolean value indicating whether this
47         Portal System is part of a Portal consisting of three
48         Portal Systems or not. Value 1 stands for a Portal of
49         three Portal Systems, value 0 stands for a Portal of
50         two Portal Systems. The default value is 0."
51     REFERENCE
52         "7.4.1.1.6"
53     ::= { dot3adDrniEntry 6 }
54
```

```
1 dot3adDrniPortalSystemNumber OBJECT-TYPE
2     SYNTAX      Integer32 (1..3)
3     MAX-ACCESS  read-write
4     STATUS      current
5     DESCRIPTION
6         "A read-write identifier of this particular Portal System
7         within a Portal. It is the responsibility of the network
8         administrator to ensure that these numbers are unique among
9         the Portal Systems with the same aDrniPortalAddr (7.4.1.1.4)."
```

```
10    REFERENCE
11        "7.4.1.1.7"
12    ::= { dot3adDrniEntry 7 }
```

```
13
14 dot3adDrniIntraPortalLinkList OBJECT-TYPE
15     SYNTAX      PortalLinkList
16     MAX-ACCESS  read-write
17     STATUS      current
18     DESCRIPTION
19         "Read-write list of the Interface Identifiers of the Ports to
20         the Intra-Portal Links assigned to this Distributed Relay. Each
21         Interface Identifier, a Port ID (6.3.4), has the two least
22         significant bits of its Port Priority (7.3.2.1.15) configured
23         to match the Portal System Number of the attached Portal
24         System. The number of IPLs in the list depends on the Portal
25         topology. For a Portal of three Portal Systems two or three
26         IPLs can be used, for a Portal of two Portal Systems a single
27         IPL is required and for a single Portal System no IPL is
28         required."
```

```
29    REFERENCE
30        "7.4.1.1.8"
31    ::= { dot3adDrniEntry 8 }
```

```
32
33 dot3adDrniAggregator OBJECT-TYPE
34     SYNTAX      InterfaceIndex
35     MAX-ACCESS  read-write
36     STATUS      current
37     DESCRIPTION
38         "An INTEGER that matches the syntax of an Interface Identifier.
39         Read-write Interface Identifier of the Aggregator Port assigned
40         to this Distributed Relay."
```

```
41    REFERENCE
42        "7.4.1.1.9"
43    ::= { dot3adDrniEntry 9 }
```

```
44
45
46 dot3adDrniNeighborAdminConvGatewayListDigest OBJECT-TYPE
47     SYNTAX      OCTET STRING (SIZE (16))
48     MAX-ACCESS  read-write
49     STATUS      current
50     DESCRIPTION
51         "The value for the digest of the prioritized Gateway
52         Conversation ID-to-Gateway assignments of the Neighbor
53         Portal System, assigned by administrator or System
54         policy for use when the Neighbor Portal System's
```

```
1         information is unknown. Its default value is set to NULL."
2     REFERENCE
3         "7.4.1.1.11"
4     DEFVAL { 'H' }
5     ::= { dot3adDrniEntry 10 }
6
7
8     dot3adDrniNeighborAdminConvPortListDigest OBJECT-TYPE
9     SYNTAX      OCTET STRING (SIZE (16))
10    MAX-ACCESS  read-write
11    STATUS      current
12    DESCRIPTION
13        "The value for the digest of the prioritized Port Conversation
14        ID-to-Aggregation Port assignments of the Neighbor Portal
15        System, assigned by administrator or System policy for use when
16        the Neighbor Portal System's information is unknown. Its default
17        value is set to NULL."
18    REFERENCE
19        "7.4.1.1.12"
20    DEFVAL { 'H' }
21    ::= { dot3adDrniEntry 11 }
22
23    dot3adDrniGatewayAlgorithm OBJECT-TYPE
24    SYNTAX      OCTET STRING (SIZE (4))
25    MAX-ACCESS  read-write
26    STATUS      current
27    DESCRIPTION
28        "This object identifies the algorithm used by the DR Function
29        to assign frames to a Gateway Conversation ID. Table 9-7
30        provides the IEEE 802.1 OUI (00-80-C2) Gateway Algorithm
31        encodings. A SEQUENCE OF OCTETS consisting of an OUI or CID
32        and one following octet."
33    REFERENCE
34        "7.4.1.1.13"
35    ::= { dot3adDrniEntry 12 }
36
37    dot3adDrniNeighborAdminGatewayAlgorithm OBJECT-TYPE
38    SYNTAX      OCTET STRING (SIZE (16))
39    MAX-ACCESS  read-write
40    STATUS      current
41    DESCRIPTION
42        "This object identifies the value for the Gateway algorithm
43        of the Neighbor Portal System, assigned by administrator or
44        System policy for use when the Neighbor Portal System's
45        information is unknown. Table 9-7 provides the IEEE 802.1
46        OUI (00-80-C2) Gateway Algorithm encodings. Its default value
47        is set to NULL. A SEQUENCE OF OCTETS consisting of a three-octet
48        OUI or CID and one following octet."
49    REFERENCE
50        "7.4.1.1.14"
51    DEFVAL { 'H' }
52    ::= { dot3adDrniEntry 13 }
53
54    dot3adDrniNeighborAdminPortAlgorithm OBJECT-TYPE
```

```
1      SYNTAX      OCTET STRING (SIZE (16))
2      MAX-ACCESS  read-write
3      STATUS      current
4      DESCRIPTION
5          "This object identifies the value for the Port Algorithm of
6          the Neighbor Portal System, assigned by administrator or
7          System policy for use when the Neighbor Portal System's
8          information is unknown. Table 6-4 provides the IEEE 802.1
9          OUI (00-80-C2) Port Algorithm encodings. Its default value
10         is set to NULL. A SEQUENCE OF OCTETS consisting of a
11         three-octet OUI or CID and one following octet."
12     REFERENCE
13         "7.4.1.1.15"
14     DEFVAL { 'H' }
15     ::= { dot3adDrniEntry 14 }
16
17 dot3adDrniNeighborAdminDRCPState OBJECT-TYPE
18     SYNTAX      DrcpState
19     MAX-ACCESS  read-write
20     STATUS      current
21     DESCRIPTION
22         "A string of 8 bits, corresponding to the administrative
23         values of DRCP_State [item s] in 9.4.3.2] as transmitted by
24         this Portal System in DRCPDUs. The first bit corresponds to
25         bit 0 of DRCP_State (HomeGateway), the second bit corresponds
26         to bit 1 (NeighborGateway), the third bit corresponds to bit 2
27         (OtherGateway), the fourth bit corresponds to bit 3
28         (IppActivity), the fifth bit corresponds to bit 4 (Timeout),
29         the sixth bit corresponds to bit 5 (GatewaySync), the seventh
30         bit corresponds to bit 6 (PortSync), and the eighth bit
31         corresponds to bit 7 (Expired). These values allow
32         administrative control over the values of HomeGateway,
33         NeighborGateway, OtherGateway, IppActivity, and Timeout. Their
34         values are by default set to FALSE. This attribute value is
35         read-write."
36     REFERENCE
37         "7.4.1.1.16"
38     ::= { dot3adDrniEntry 15 }
39
40 dot3adDrniEncapsulationMethod OBJECT-TYPE
41     SYNTAX      OCTET STRING (SIZE (4))
42     MAX-ACCESS  read-write
43     STATUS      current
44     DESCRIPTION
45         "This managed object is applicable only when Network / IPL
46         sharing by time (9.3.2.1) or Network / IPL sharing by tag
47         (9.3.2.2) or Network / IPL sharing by encapsulation (9.3.2.3)
48         is supported. This object identifies the value representing the
49         encapsulation method that is used to transport IPL frames to
50         the Neighbor Portal System when the IPL and network link are
51         sharing the same physical link. It consists of the three-octet
52         OUI or CID identifying the organization which is responsible
53         for this encapsulation and one following octet used to identify
54         the encapsulation method defined by that organization.
```

1 Table 9-10 provides the IEEE 802.1 OUI (00-80-C2) encapsulation
2 method encodings. A Default value of 0x00-80-C2-00 indicates
3 that the IPL is using a separate physical or Aggregation link.
4 A value of 1 indicates that Network / IPL sharing by time
5 (9.3.2.1) is used. A value of 2 indicates that the encapsulation
6 method used is the same as the one used by network frames and
7 that Network / IPL sharing by tag (9.3.2.2) is used. A SEQUENCE
8 OF OCTETS consisting of OUI or CID and one following octet."

9 REFERENCE

10 "7.4.1.1.17"

11 DEFVAL { '0080C200'H }

12 ::= { dot3adDrniEntry 16 }

13
14 dot3adDrniDRPortConversationPasses OBJECT-TYPE

15 SYNTAX OCTET STRING (SIZE (512))

16 MAX-ACCESS read-only

17 STATUS current

18 DESCRIPTION

19 "A read-only current operational vector of Boolean values, with
20 one value for each possible Port Conversation ID. A 1 indicates
21 that the Port Conversation ID is allowed to be distributed
22 through this DR Function's Aggregator, and a 0 indicates that
23 it cannot. aDrniDRPortConversationPasses is referencing the
24 current value of Drni_Portal_System_Port_Conversation
25 (9.3.4.2)."

26 REFERENCE

27 "7.4.1.1.20"

28 ::= { dot3adDrniEntry 17 }

29
30 dot3adDrniDRGatewayConversationPasses OBJECT-TYPE

31 SYNTAX OCTET STRING (SIZE (512))

32 MAX-ACCESS read-only

33 STATUS current

34 DESCRIPTION

35 "A read-only current operational vector of Boolean values,
36 with one value for each possible Gateway Conversation ID. A
37 1 indicates that the Gateway Conversation ID is allowed to pass
38 through this DR Function's Gateway, and a 0 indicates that it
39 cannot. aDrniDRGatewayConversationPasses is referencing the
40 current value of Drni_Portal_System_Gateway_Conversation
41 (9.3.4.2)."

42 REFERENCE

43 "7.4.1.1.21"

44 ::= { dot3adDrniEntry 18 }

45
46 dot3adDrniPSI OBJECT-TYPE

47 SYNTAX TruthValue

48 MAX-ACCESS read-only

49 STATUS current

50 DESCRIPTION

51 "A read-only Boolean value providing the value of PSI, which
52 indicates whether this Portal System is isolated from the
53 other Portal Systems within the same Portal ('TRUE') or not
54 ('FALSE')."

```

1      REFERENCE
2          "7.4.1.1.22"
3      ::= { dot3adDrniEntry 19 }
4
5      dot3adDrniPortConversationControl OBJECT-TYPE
6          SYNTAX      TruthValue
7          MAX-ACCESS  read-write
8          STATUS      current
9          DESCRIPTION
10             "A read-write Boolean value that controls the operation of the
11             updatedDRFHomeState (9.4.11). When set to 'TRUE' the Home
12 Gateway
13             Vector is set equal to Drni_Portal_System_Port_Conversation.
14             Setting this object to 'TRUE' is only possible when the
15 Gateway
16             algorithm and the Port algorithm use the same distributions
17             methods. The default is 'FALSE', indicating that the Home Gateway
18             Vector is controlled by the network control protocol."
19      REFERENCE
20          "7.4.1.1.23"
21      DEFVAL { false }
22      ::= { dot3adDrniEntry 20 }
23
24
25      dot3adDrniIntraPortalPortProtocolDA OBJECT-TYPE
26          SYNTAX      MacAddress
27          MAX-ACCESS  read-write
28          STATUS      current
29          DESCRIPTION
30             "A 6-octet read-write MAC Address value specifying the
31             destination address to be used when sending DRCPDUs,
32             corresponding to the value of DRCP_Protocol_DA in 9.4.4.1.3.
33             Its values is one of the addresses selected from Table 9-4
34             and its default shall be the IEEE 802.1 Nearest non-TPMR
35             Bridge group address (01-80-C2-00-00-03)."

```

1 Conversation ID. This selection priority list, a sequence of
2 integers for each Gateway Conversation ID, is a list of Portal
3 System Numbers in the order of preference, highest to lowest,
4 for the corresponding preferred Portal System's Gateway to
5 carry that Conversation. NOTE - To the extent that the network
6 administrator fails to configure the same values for the
7 aDrniConvAdminGateway[] variables in all of the DR Functions
8 of a Portal, frames can be misdirected. The Distributed Relay
9 Control Protocol (DRCP, 9.4) detects such misconfiguration."

10 REFERENCE

11 "7.4.1.1.10"
12 ::= { dot3adDrni 2 }

13
14

15 dot3adDrniConvAdminGatewayEntry OBJECT-TYPE

16 SYNTAX Dot3adDrniConvAdminGatewayEntry

17 MAX-ACCESS not-accessible

18 STATUS current

19 DESCRIPTION

20 "A Gateway selection priority list for the Distributed Relay
21 for the referenced Gateway Conversation ID. This selection
22 priority list, a sequence of integers for each Gateway
23 Conversation ID, is a list of Portal System Numbers in the
24 order of preference, highest to lowest, for the corresponding
25 preferred Portal System's Gateway to carry that Conversation."

26 INDEX { dot3adDrniGatewayConversationID, dot3adDrniIndex }
27 ::= { dot3adDrniConvAdminGatewayTable 1 }

28
29

29 Dot3adDrniConvAdminGatewayEntry ::=

30 SEQUENCE {
31 dot3adDrniGatewayConversationID
32 Integer32,
33 dot3adDrniConvAdminGatewayList
34 DrniConvAdminGatewayList
35 }

36
37

37 dot3adDrniGatewayConversationID OBJECT-TYPE

38 SYNTAX Integer32 (0..4095)

39 MAX-ACCESS not-accessible

40 STATUS current

41 DESCRIPTION

42 "An identifier for a Gateway Conversation."

43 ::= { dot3adDrniConvAdminGatewayEntry 1 }

44
45

45 dot3adDrniConvAdminGatewayList OBJECT-TYPE

46 SYNTAX DrniConvAdminGatewayList

47 MAX-ACCESS read-write

48 STATUS current

49 DESCRIPTION

50 "Priority list of Portal System Numbers in order of preference
51 from highest to lowest."

52 REFERENCE

53 "7.4.1.1.10"
54 ::= { dot3adDrniConvAdminGatewayEntry 2 }

```

1
2  -----
3  -- DRNI IPL Encap Map Table
4  -----
5
6  dot3adDrniIPLEncapMapTable OBJECT-TYPE
7      SYNTAX      SEQUENCE OF Dot3adDrniIPLEncapMapEntry
8      MAX-ACCESS  not-accessible
9      STATUS      current
10     DESCRIPTION
11         "This managed object is applicable only when Network / IPL
12         sharing by tag (9.3.2.2) or Network / IPL sharing by
13     encapsulation
14         (9.3.2.3) is supported. Each entry represents the value of the
15         identifier used for an IPL frame associated with that Gateway
16         Conversation ID for the encapsulation method specified in
17         7.4.1.1.17."
18     REFERENCE
19         "7.4.1.1.18"
20     ::= { dot3adDrni 3 }
21
22     dot3adDrniIPLEncapMapEntry OBJECT-TYPE
23         SYNTAX      Dot3adDrniIPLEncapMapEntry
24         MAX-ACCESS  not-accessible
25         STATUS      current
26         DESCRIPTION
27             "An entry represents the value of the identifier used for an
28             IPL frame associated with that Gateway Conversation ID for the
29             encapsulation method specified in 7.4.1.1.17."
30         INDEX      { dot3adDrniGatewayConversationID, dot3adDrniIndex }
31         ::= { dot3adDrniIPLEncapMapTable 1 }
32
33     Dot3adDrniIPLEncapMapEntry ::=
34         SEQUENCE {
35             dot3adDrniIPLFrameIdValue
36                 Integer32
37         }
38
39     dot3adDrniIPLFrameIdValue OBJECT-TYPE
40         SYNTAX      Integer32
41         MAX-ACCESS  read-write
42         STATUS      current
43         DESCRIPTION
44             "The value of the identifier used for an IPL frame associated
45             with that Gateway Conversation ID for the encapsulation
46             method."
47         REFERENCE
48             "7.4.1.1.18"
49         ::= { dot3adDrniIPLEncapMapEntry 2 }
50
51  -----
52  -- DRNI Net Encap Map Table
53  -----
54

```



```

1 dot3adDrniNetEncapMapTable OBJECT-TYPE
2     SYNTAX      SEQUENCE OF Dot3adDrniNetEncapMapEntry
3     MAX-ACCESS  not-accessible
4     STATUS      current
5     DESCRIPTION
6         "This managed object is applicable only when Network /
7         IPL sharing by tag (9.3.2.2) is supported. Each entry
8         represents the translated value of the identifier used
9         for a network frame associated with that Gateway
10        Conversation ID when the method specified in 7.4.1.1.17
11        is the Network / IPL sharing by tag method specified in
12        9.3.2.2 and the network frames need to share the tag
13        space used by IPL frames."
14    REFERENCE
15        "7.4.1.1.19"
16    ::= { dot3adDrni 4 }
17
18 dot3adDrniNetEncapMapEntry OBJECT-TYPE
19     SYNTAX      Dot3adDrniNetEncapMapEntry
20     MAX-ACCESS  not-accessible
21     STATUS      current
22     DESCRIPTION
23         "An entry represents the translated value of the identifier
24         used for a network frame associated with that Gateway
25         Conversation ID when the method specified in 7.4.1.1.12 is the
26         Network / IPL sharing by tag method specified in 9.3.2.2 and
27         the network frames need to share the tag space used by IPL
28         frames."
29     INDEX       { dot3adDrniGatewayConversationID, dot3adDrniIndex }
30     ::= { dot3adDrniNetEncapMapTable 1 }
31
32 Dot3adDrniNetEncapMapEntry ::=
33     SEQUENCE {
34         dot3adDrniNetFrameIdValue
35         Integer32
36     }
37
38 dot3adDrniNetFrameIdValue OBJECT-TYPE
39     SYNTAX      Integer32
40     MAX-ACCESS  read-write
41     STATUS      current
42     DESCRIPTION
43         "The translated value of the identifier used for a network
44         frame associated that Gateway Conversation ID."
45     REFERENCE
46         "7.4.1.1.19"
47     ::= { dot3adDrniNetEncapMapEntry 1 }
48
49     -----
50     -- IPP Attribute Table
51     -----
52
53 dot3adIPPAtributeTable OBJECT-TYPE
54     SYNTAX      SEQUENCE OF Dot3adIPPAtributeEntry

```

```
1      MAX-ACCESS not-accessible
2      STATUS current
3      DESCRIPTION
4          "A table that contains information about every
5          IPP that is associated with this System."
6      REFERENCE
7          "7.4.2"
8      ::= { dot3adIPP 1 }
9
10     dot3adIPPAttributeEntry OBJECT-TYPE
11         SYNTAX Dot3adIPPAttributeEntry
12         MAX-ACCESS not-accessible
13         STATUS current
14         DESCRIPTION
15             "An entry containing Attributes for an IPP."
16         INDEX { dot3adIPPIndex }
17         ::= { dot3adIPPAttributeTable 1 }
18
19     Dot3adIPPAttributeEntry ::=
20         SEQUENCE {
21             dot3adIPPIndex
22                 InterfaceIndex,
23             dot3adIPPPortConversationPasses
24                 OCTET STRING,
25             dot3adIPPGatewayConversationDirection
26                 OCTET STRING,
27             dot3adIPPAdminState
28                 AggState,
29             dot3adIPPOperState
30                 AggState,
31             dot3adIPPTimeOfLastOperChange
32                 Integer32
33         }
34
35     dot3adIPPIndex OBJECT-TYPE
36         SYNTAX InterfaceIndex
37         MAX-ACCESS not-accessible
38         STATUS current
39         DESCRIPTION
40             "The unique identifier allocated to this IPP by the
41             local Portal System. This attribute identifies an
42             IPP instance among the subordinate managed objects of
43             the containing object. This value is read-only.
44             NOTE-The aIPPID is represented in the SMIV2 MIB as
45             an ifIndex-see D.5."
46         REFERENCE
47             "7.4.2.1.1"
48         ::= { dot3adIPPAttributeEntry 1 }
49
50     dot3adIPPPortConversationPasses OBJECT-TYPE
51         SYNTAX OCTET STRING (SIZE (512))
52         MAX-ACCESS read-only
53         STATUS current
54         DESCRIPTION
```

```
1         "A read-only current operational vector of Boolean
2         values, with one value for each possible Port
3         Conversation ID. A 1 indicates that the Port Conversation
4         ID is allowed to be transmitted through this IPP, and a 0
5         indicates that it cannot. aDrniIPPPortConversationPasses is
6         referencing the current value of Ipp_Port_Conversation_Passes
7         (9.3.4.3)."
```

8 REFERENCE

```
9         "7.4.2.1.2"
10        ::= { dot3adIPPAttributeEntry 2 }
```

11

12 dot3adIPPGatewayConversationDirection OBJECT-TYPE

```
13     SYNTAX      OCTET STRING (SIZE (512))
14     MAX-ACCESS  read-only
15     STATUS      current
16     DESCRIPTION
17         "A read-only current operational vector of Boolean values,
18         with one value for each possible Gateway Conversation ID.
19         A 1 indicates that the Gateway Conversation ID is assigned to
20         Gateways reachable through this IPP, and a 0 indicates that
21         the Gateway for the indexed Gateway Conversation ID is not
22         reachable through this IPP.
23         aDrniIPPGatewayConversationDirection is referencing the
24         current value of Ipp_Gateway_Conversation_Direction (9.3.4.3)."
```

25 REFERENCE

```
26         "7.4.2.1.3"
27        ::= { dot3adIPPAttributeEntry 3 }
```

28

29 dot3adIPPAdminState OBJECT-TYPE

```
30     SYNTAX      AggState
31     MAX-ACCESS  read-write
32     STATUS      current
33     DESCRIPTION
34         "This read-write value defines the administrative state of
35         the IPP. A value of 'up' indicates that the operational state
36         of the IPP (aIPPOperState) is permitted to be either up or
37         down. A value of 'down' forces the operational state of the IPP
38         to be down. A GET operation returns the current administrative
39         state. A SET operation changes the administrative state to a
40         new value."
```

41 REFERENCE

```
42         "7.4.2.1.4"
43        ::= { dot3adIPPAttributeEntry 4 }
```

44

45 dot3adIPPOperState OBJECT-TYPE

```
46     SYNTAX      AggState
47     MAX-ACCESS  read-only
48     STATUS      current
49     DESCRIPTION
50         "This read-only value defines the operational state of the
51         IPP. The operational state is 'up' if the IPL is operational,
52         and if the value of aIPPAdminState for the IPP is also 'up'.
53         If the IPL is not operational, or if the administrative state
54         of the IPP (aIPPAdminState) is 'down', then the operational
```

1 state is 'down.' An operational state of 'up' indicates that
2 the IPP is available for use by the DR Function; a value of
3 'down' indicates that the IPP is not available for use by the
4 DR Function."

5 REFERENCE
6 "7.4.2.1.5"
7 ::= { dot3adIPPAttributeEntry 5 }

8
9 dot3adIPPTimeOfLastOperChange OBJECT-TYPE

10 SYNTAX Integer32
11 MAX-ACCESS read-only
12 STATUS current
13 DESCRIPTION
14 "The time at which the interface entered its current
15 operational state, in terms of centiseconds since the
16 system was last reset. If the current state was entered
17 prior to the last reinitialization of the local network
18 management subsystem, then this object contains a value of
19 zero. The ifLastChange object in the Interfaces MIB defined
20 in IETF RFC 2863 is a suitable object for supplying a
21 value for aIPPTimeOfLastOperChange. This value is read-only."

22 REFERENCE
23 "7.4.2.1.6"
24 ::= { dot3adIPPAttributeEntry 6 }

25
26
27 -----
28 -- IPP Statistics Table
29 -----

30
31 dot3adIPPStatsTable OBJECT-TYPE

32 SYNTAX SEQUENCE OF Dot3adIPPStatsEntry
33 MAX-ACCESS not-accessible
34 STATUS current
35 DESCRIPTION
36 "A table that contains information for IPP
37 statistics. A row appears in this table for
38 each IPP in the system."

39 REFERENCE
40 "7.4.3"
41 ::= { dot3adIPP 2 }

42
43 dot3adIPPStatsEntry OBJECT-TYPE

44 SYNTAX Dot3adIPPStatsEntry
45 MAX-ACCESS not-accessible
46 STATUS current
47 DESCRIPTION
48 "An entry containing Statistics for an IPP."
49 INDEX { dot3adIPPIndex }
50 ::= { dot3adIPPStatsTable 1 }

51
52 Dot3adIPPStatsEntry ::=

53 SEQUENCE {
54 dot3adIPPStatsDRCPDUsRx

```

1         Counter64,
2         dot3adIPPStatsIllegalRx
3         Counter64,
4         dot3adIPPStatsDRCPDUsTx
5         Counter64
6     }
7
8 dot3adIPPStatsDRCPDUsRx OBJECT-TYPE
9     SYNTAX      Counter64
10    UNITS
11        "frames"
12    MAX-ACCESS  read-only
13    STATUS      current
14    DESCRIPTION
15        "The number of valid DRCPDUs received on this IPP. This value
16        is read-only."
17    REFERENCE
18        "7.4.3.1.2"
19    ::= { dot3adIPPStatsEntry 1 }
20
21 dot3adIPPStatsIllegalRx OBJECT-TYPE
22    SYNTAX      Counter64
23    UNITS
24        "frames"
25    MAX-ACCESS  read-only
26    STATUS      current
27    DESCRIPTION
28        "The number of frames received that carry the DRCP Ethernet
29        Type value (9.4.2.4), but contain a badly formed PDU. This
30        value is read-only."
31    REFERENCE
32        "7.4.3.1.3"
33    ::= { dot3adIPPStatsEntry 2 }
34
35
36 dot3adIPPStatsDRCPDUsTx OBJECT-TYPE
37    SYNTAX      Counter64
38    UNITS
39        "frames"
40    MAX-ACCESS  read-only
41    STATUS      current
42    DESCRIPTION
43        "The number of DRCPDUs transmitted on this IPP. This value
44        is read-only."
45    REFERENCE
46        "7.4.3.1.4"
47    ::= { dot3adIPPStatsEntry 3 }
48
49
50 -----
51 -- IPP Debug Table
52 -----
53
54 dot3adIPPDebugTable OBJECT-TYPE

```

```

1      SYNTAX      SEQUENCE OF Dot3adIPPDebugEntry
2      MAX-ACCESS  not-accessible
3      STATUS      current
4      DESCRIPTION
5          "A table that contains IPP debug information.
6          A row appears in this table for each IPP in
7          the system."
8      REFERENCE
9          "7.4.4"
10     ::= { dot3adIPP 3 }
11
12     dot3adIPPDebugEntry OBJECT-TYPE
13     SYNTAX      Dot3adIPPDebugEntry
14     MAX-ACCESS  not-accessible
15     STATUS      current
16     DESCRIPTION
17         "An entry containing Debug Information for an IPP."
18     INDEX { dot3adIPPIndex }
19     ::= { dot3adIPPDebugTable 1 }
20
21     Dot3adIPPDebugEntry ::=
22     SEQUENCE {
23         dot3adIPPDebugDRCPRxState
24             Integer32,
25         dot3adIPPDebugLastRxTime
26             TimeTicks,
27         dot3adIPPDebugDifferPortalReason
28             SnmpAdminString
29     }
30
31     dot3adIPPDebugDRCPRxState OBJECT-TYPE
32     SYNTAX      INTEGER {
33                 current(1),
34                 expired(2),
35                 defaulted(3),
36                 initialize(4),
37                 reportToManagement(5)
38             }
39     MAX-ACCESS  read-only
40     STATUS      current
41     DESCRIPTION
42         "This attribute holds the value 'current' if the DRCPDU
43         Receive state machine for the IPP is in the CURRENT state,
44         'expired' if the DRCPDU Receive state machine is in the
45         EXPIRED state, 'defaulted' if the DRCPDU Receive state machine
46         is in the DEFAULTED state, 'initialize' if the DRCPDU Receive
47         state machine is in the INITIALIZE state, or
48         'reportToManagement' if the Receive state machine is in the
49         REPORT_TO_MANAGEMENT state. This value is read-only."
50     REFERENCE
51         "7.4.4.1.2"
52     ::= { dot3adIPPDebugEntry 1 }
53
54

```

```

1 dot3adIPPDebugLastRxTime OBJECT-TYPE
2     SYNTAX          TimeTicks
3     MAX-ACCESS      read-only
4     STATUS          current
5     DESCRIPTION
6         "The time at which the last DRCPDU was received by this IPP,
7         in terms of centiseconds since the system was last reset.
8         The ifLastChange object in the Interfaces MIB defined in IETF
9         RFC 2863 is a suitable object for supplying a value for
10        aDrniIPPDebugLastRxTime. This value is read-only."
11    REFERENCE
12        "7.4.4.1.3"
13    ::= { dot3adIPPDebugEntry 2 }
14
15 dot3adIPPDebugDifferPortalReason OBJECT-TYPE
16     SYNTAX          SnmpAdminString
17     MAX-ACCESS      read-only
18     STATUS          current
19     DESCRIPTION
20        "A human-readable text string indicating the most recent set
21        of variables that are responsible for setting the variable
22        Differ_Portal or Differ_Conf_Portal (9.4.8) on this IPP to
23        TRUE. This value is read-only."
24    REFERENCE
25        "7.4.4.1.4"
26    ::= { dot3adIPPDebugEntry 3 }
27
28
29    -----
30    -- Aggregator Notifications
31    -----
32
33 dot3adAggLinkUpNotification NOTIFICATION-TYPE
34     -- OBJECTS { }
35     STATUS          current
36     DESCRIPTION
37        "When aAggLinkUpDownNotificationEnable is set to
38        'enabled,' a Link Up notification is generated when
39        the Operational State of the Aggregator changes from
40        'down' to 'up.' When aAggLinkUpDownNotificationEnable
41        is set to 'disabled,' no Link Up notifications are
42        generated. The notification carries the identifier of
43        the Aggregator whose state has changed."
44     ::= { lagMIBNotifications 1 }
45
46 dot3adAggLinkDownNotification NOTIFICATION-TYPE
47     STATUS          current
48     DESCRIPTION
49        "When aAggLinkUpDownNotificationEnable is set to
50        'enabled,' a Link Down notification is generated when
51        the Operational State of the Aggregator changes from
52        'up' to 'down.' When aAggLinkUpDownNotificationEnable
53        is set to 'disabled,' no Link Down notifications are
54        generated. The notification carries the identifier of

```

```

1         the Aggregator whose state has changed."
2     ::= { lagMIBNotifications 2 }
3
4
5     -----
6     -- IEEE 802.3ad MIB - Conformance Information
7     -----
8
9     dot3adAggGroups OBJECT IDENTIFIER
10    ::= { dot3adAggConformance 1 }
11
12    dot3adAggCompliances OBJECT IDENTIFIER
13    ::= { dot3adAggConformance 2 }
14
15    -----
16    -- units of conformance
17    -----
18
19    dot3adAggGroup OBJECT-GROUP
20    OBJECTS {
21        dot3adAggActorSystemID,
22        dot3adAggActorSystemPriority,
23        dot3adAggAggregateOrIndividual,
24        dot3adAggActorAdminKey,
25        dot3adAggMACAddress,
26        dot3adAggActorOperKey,
27        dot3adAggPartnerSystemID,
28        dot3adAggPartnerSystemPriority,
29        dot3adAggPartnerOperKey,
30        dot3adAggCollectorMaxDelay
31    }
32    STATUS          current
33    DESCRIPTION
34        "A collection of objects providing information about an
35        aggregation."
36    ::= { dot3adAggGroups 1 }
37
38    dot3adAggPortListGroup OBJECT-GROUP
39    OBJECTS {
40        dot3adAggPortListGroupPorts
41    }
42    STATUS          current
43    DESCRIPTION
44        "A collection of objects providing information about every
45        port in an aggregation."
46    ::= { dot3adAggGroups 2 }
47
48    dot3adAggPortGroup OBJECT-GROUP
49    OBJECTS {
50        dot3adAggPortActorSystemPriority,
51        dot3adAggPortActorSystemID,
52        dot3adAggPortActorAdminKey,
53        dot3adAggPortActorOperKey,
54        dot3adAggPortPartnerAdminSystemPriority,

```



```

1      dot3adAggPortPartnerOperSystemPriority,
2      dot3adAggPortPartnerAdminSystemID,
3      dot3adAggPortPartnerOperSystemID,
4      dot3adAggPortPartnerAdminKey,
5      dot3adAggPortPartnerOperKey,
6      dot3adAggPortSelectedAggID,
7      dot3adAggPortAttachedAggID,
8      dot3adAggPortActorPort,
9      dot3adAggPortActorPortPriority,
10     dot3adAggPortPartnerAdminPort,
11     dot3adAggPortPartnerOperPort,
12     dot3adAggPortPartnerAdminPortPriority,
13     dot3adAggPortPartnerOperPortPriority,
14     dot3adAggPortActorAdminState,
15     dot3adAggPortActorOperState,
16     dot3adAggPortPartnerAdminState,
17     dot3adAggPortPartnerOperState,
18     dot3adAggPortAggregateOrIndividual
19   }
20   STATUS      current
21   DESCRIPTION
22     "A collection of objects providing information about every
23     port in an aggregation."
24   ::= { dot3adAggGroups 3 }
25
26
27   dot3adAggPortStatsGroup OBJECT-GROUP
28     OBJECTS {
29       dot3adAggPortStatsLACPDUrx,
30       dot3adAggPortStatsMarkerPDURx,
31       dot3adAggPortStatsMarkerResponsePDURx,
32       dot3adAggPortStatsUnknownRx,
33       dot3adAggPortStatsIllegalRx,
34       dot3adAggPortStatsLACPDUtx,
35       dot3adAggPortStatsMarkerPDUTx,
36       dot3adAggPortStatsMarkerResponsePDUTx
37     }
38     STATUS      current
39     DESCRIPTION
40       "A collection of objects providing information about every
41       port in an aggregation."
42     ::= { dot3adAggGroups 4 }
43
44
45   dot3adAggPortDebugGroup OBJECT-GROUP
46     OBJECTS {
47       dot3adAggPortDebugRxState,
48       dot3adAggPortDebugLastRxTime,
49       dot3adAggPortDebugMuxState,
50       dot3adAggPortDebugMuxReason,
51       dot3adAggPortDebugActorChurnState,
52       dot3adAggPortDebugPartnerChurnState,
53       dot3adAggPortDebugActorChurnCount,
54       dot3adAggPortDebugPartnerChurnCount,

```

```
1         dot3adAggPortDebugActorSyncTransitionCount,
2         dot3adAggPortDebugPartnerSyncTransitionCount,
3         dot3adAggPortDebugActorChangeCount,
4         dot3adAggPortDebugPartnerChangeCount
5     }
6     STATUS      current
7     DESCRIPTION
8         "A collection of objects providing debug information about
9         every aggregated port."
10    ::= { dot3adAggGroups 5 }
11
12    dot3adTablesLastChangedGroup OBJECT-GROUP
13    OBJECTS {
14        dot3adTablesLastChanged
15    }
16    STATUS      current
17    DESCRIPTION
18        "A collection of objects providing information about the time
19        of changes to the configuration of aggregations and their
20        ports."
21    ::= { dot3adAggGroups 6 }
22
23    dot3adAggPortProtocolDAGroup OBJECT-GROUP
24    OBJECTS {
25        dot3adAggPortProtocolDA
26    }
27    STATUS      current
28    DESCRIPTION
29        "A collection of objects providing information about the
30        protocol destination address in use for ports in an
31        aggregation."
32    ::= { dot3adAggGroups 7 }
33
34    dot3adAggNotificationGroup NOTIFICATION-GROUP
35    NOTIFICATIONS {
36        dot3adAggLinkUpNotification,
37        dot3adAggLinkDownNotification
38    }
39    STATUS      current
40    DESCRIPTION
41        "A collection of notifications providing information about
42        the aggregation."
43    ::= { dot3adAggGroups 8 }
44
45    dot3adAggXGroup OBJECT-GROUP
46    OBJECTS {
47        dot3adAggDescription,
48        dot3adAggName,
49        dot3adAggAdminState,
50        dot3adAggOperState,
51        dot3adAggTimeOfLastOperChange,
52        dot3adAggDataRate,
53        dot3adAggFramesTxOK,
54        dot3adAggFramesRxOK,
```

```
1         dot3adAggLinkUpDownNotificationEnable
2     }
3     STATUS    current
4     DESCRIPTION
5         "A collection of extension entries providing information on an
6         aggregation."
7     ::= { dot3adAggGroups 9 }
8
9     dot3adAggRecommendedGroup OBJECT-GROUP
10    OBJECTS {
11        dot3adAggOctetsTxOK,
12        dot3adAggOctetsRxOK,
13        dot3adAggFramesDiscardedOnTx,
14        dot3adAggFramesDiscardedOnRx,
15        dot3adAggFramesWithTxErrors,
16        dot3adAggFramesWithRxErrors,
17        dot3adAggUnknownProtocolFrames
18    }
19    STATUS    current
20    DESCRIPTION
21        "A collection of recommended objects providing information
22        about an aggregation."
23    ::= { dot3adAggGroups 10 }
24
25    dot3adAggOptionalGroup OBJECT-GROUP
26    OBJECTS {
27        dot3adAggMulticastFramesTxOK,
28        dot3adAggMulticastFramesRxOK,
29        dot3adAggBroadcastFramesTxOK,
30        dot3adAggBroadcastFramesRxOK
31    }
32    STATUS    current
33    DESCRIPTION
34        "A collection of optional objects providing information about
35        an aggregation."
36    ::= { dot3adAggGroups 11 }
37
38    dot3adPerServiceFrameDistGroup OBJECT-GROUP
39    OBJECTS {
40        dot3adAggConversationAdminLinkList,
41        dot3adAggPortAlgorithm,
42        dot3adAggPartnerAdminPortAlgorithm,
43        dot3adAggPartnerAdminPortConversationListDigest,
44        dot3adAggAdminDiscardWrongConversation,
45        dot3adAggPartnerAdminConvServiceMappingDigest,
46        dot3adAggAdminServiceConversationServiceIDList,
47        dot3adAggPortLinkNumberId,
48        dot3adAggPortPartnerAdminLinkNumberId,
49        dot3adAggPortOperConversationPasses,
50        dot3adAggPortOperConversationCollected,
51        dot3adAggPortWTRTime
52    }
53    STATUS    current
54    DESCRIPTION
```

```
1           "A collection of objects providing information about
2           Per-Service Frame Distribution."
3           ::= { dot3adAggGroups 12 }
4
5 dot3adAggPortDebugXGroup OBJECT-GROUP
6     OBJECTS {
7         dot3adAggPortDebugActorCDSChurnState,
8         dot3adAggPortDebugPartnerCDSChurnState,
9         dot3adAggPortDebugActorCDSChurnCount,
10        dot3adAggPortDebugPartnerCDSChurnCount
11    }
12    STATUS    current
13    DESCRIPTION
14        "A collection of objects extending aggregator port debug."
15        ::= { dot3adAggGroups 13 }
16
17
18 dot3adDrniGroup OBJECT-GROUP
19     OBJECTS {
20         dot3adDrniDescription,
21         dot3adDrniName,
22         dot3adDrniPortalAddr,
23         dot3adDrniPortalPriority,
24         dot3adDrniThreePortalSystem,
25         dot3adDrniPortalSystemNumber,
26         dot3adDrniIntraPortalLinkList,
27         dot3adDrniAggregator,
28         dot3adDrniNeighborAdminConvGatewayListDigest,
29         dot3adDrniNeighborAdminConvPortListDigest,
30         dot3adDrniGatewayAlgorithm,
31         dot3adDrniNeighborAdminGatewayAlgorithm,
32         dot3adDrniNeighborAdminPortAlgorithm,
33         dot3adDrniNeighborAdminDRCPState,
34         dot3adDrniEncapsulationMethod,
35         dot3adDrniDRPortConversationPasses,
36         dot3adDrniDRGatewayConversationPasses,
37         dot3adDrniConvAdminGatewayList,
38         dot3adDrniIPLFrameIdValue,
39         dot3adDrniNetFrameIdValue,
40         dot3adDrniPSI,
41         dot3adDrniPortConversationControl,
42         dot3adDrniIntraPortalPortProtocolDA
43    }
44    STATUS    current
45    DESCRIPTION
46        "A collection of objects providing information about DRNI."
47        ::= { dot3adAggGroups 14 }
48
49 dot3adIPPGroup OBJECT-GROUP
50     OBJECTS {
51         dot3adIPPPortConversationPasses,
52         dot3adIPPGatewayConversationDirection,
53         dot3adIPPAdminState,
54         dot3adIPPOperState,
```

```

1         dot3adIPPTimeOfLastOperChange
2     }
3     STATUS    current
4     DESCRIPTION
5         "A collection of objects providing IPP information."
6     ::= { dot3adAggGroups 15 }
7
8 dot3adIPPStatsGroup  OBJECT-GROUP
9     OBJECTS {
10        dot3adIPPStatsDRCPDUsRx,
11        dot3adIPPStatsIllegalRx,
12        dot3adIPPStatsDRCPDUsTx
13    }
14    STATUS    current
15    DESCRIPTION
16        "A collection of objects providing IPP statistics information."
17    ::= { dot3adAggGroups 16 }
18
19 dot3adIPPDebugGroup  OBJECT-GROUP
20    OBJECTS {
21        dot3adIPPDebugDRCPRxState,
22        dot3adIPPDebugLastRxTime,
23        dot3adIPPDebugDifferPortalReason
24    }
25    STATUS    current
26    DESCRIPTION
27        "A collection of objects providing IPP debug information."
28    ::= { dot3adAggGroups 17 }
29
30
31    -----
32    -- compliance statements
33    -----
34
35 dot3adAggCompliance MODULE-COMPLIANCE
36    STATUS    current
37    DESCRIPTION
38        "The compliance statement for device support of
39        Link Aggregation."
40
41    MODULE
42        MANDATORY-GROUPS {
43            dot3adAggGroup,
44            dot3adAggPortGroup,
45            dot3adAggNotificationGroup,
46            dot3adTablesLastChangedGroup
47        }
48
49        GROUP        dot3adAggPortListGroup
50        DESCRIPTION
51            "This group is optional."
52
53        GROUP        dot3adAggXGroup
54        DESCRIPTION

```

```
1           "This group is optional."  
2  
3     GROUP      dot3adAggPortStatsGroup  
4     DESCRIPTION  
5           "This group is optional."  
6  
7     GROUP      dot3adAggPortDebugGroup  
8     DESCRIPTION  
9           "This group is optional."  
10  
11    GROUP      dot3adAggPortDebugXGroup  
12    DESCRIPTION  
13          "This group is optional."  
14  
15    GROUP      dot3adAggPortProtocolDAGroup  
16    DESCRIPTION  
17          "This group is optional."  
18  
19    GROUP      dot3adAggOptionalGroup  
20    DESCRIPTION  
21          "This group is optional."  
22  
23    GROUP      dot3adAggRecommendedGroup  
24    DESCRIPTION  
25          "This group is optional."  
26  
27    GROUP      dot3adPerServiceFrameDistGroup  
28    DESCRIPTION  
29          "This group is optional."  
30  
31    GROUP      dot3adDrniGroup  
32    DESCRIPTION  
33          "This group is optional."  
34  
35    GROUP      dot3adIPPGroup  
36    DESCRIPTION  
37          "This group is optional."  
38  
39          GROUP      dot3adIPPStatsGroup  
40    DESCRIPTION  
41          "This group is optional."  
42  
43    GROUP      dot3adIPPDebugGroup  
44    DESCRIPTION  
45          "This group is optional."  
46  
47 ::= { dot3adAggCompliances 1 }  
48  
49  
50 END  
51  
52  
53  
54
```

Annex E

(informative)

Distributed Bridge

Clause 9 describes the operation of DRNI in a manner agnostic to the forwarding capabilities and the specific service interface support of the Systems implementing it. In particular, the Distributed Relay in 9.2 and the DR Function in 9.3 are described in terms of generic forwarding functional elements and operations that enable support of DRNI by a wide range of Systems. A generic DR Function must support the forwarding rules that are summarized in Table 9-1 to Table 9-3 and specified by the state machine descriptions in 9.3.4 in terms of generic Gateway and Port Algorithms. The accommodation of the generic DR Function by a specific System is putting constraints to the choices of Gateway and Port Algorithms that are applicable. It is the specific forwarding relay capabilities of the supporting System (based on B-VLANs, C-VLANs, 5-tuples, etc.) that are putting limits to the Gateway Algorithms that can be supported by the DR Function in that specific Distributed System, while it is the frame distribution functions and service interface support (based on I-SIDs, S-VLANs, 5-tuples, etc.) that select the Port Algorithms that are applicable to that specific System. This Annex discusses Distributed Relays supported by IEEE Std 802.1Q Bridges.

E.1 Distributed VLAN Bridge

The simplest example involving Bridges corresponds to the case of the DRNI supported by two VLAN Bridges on a Customer Network Port (CNP) or Provider Network Port (PNP). Figure E-1 is providing the logical model of such a case.

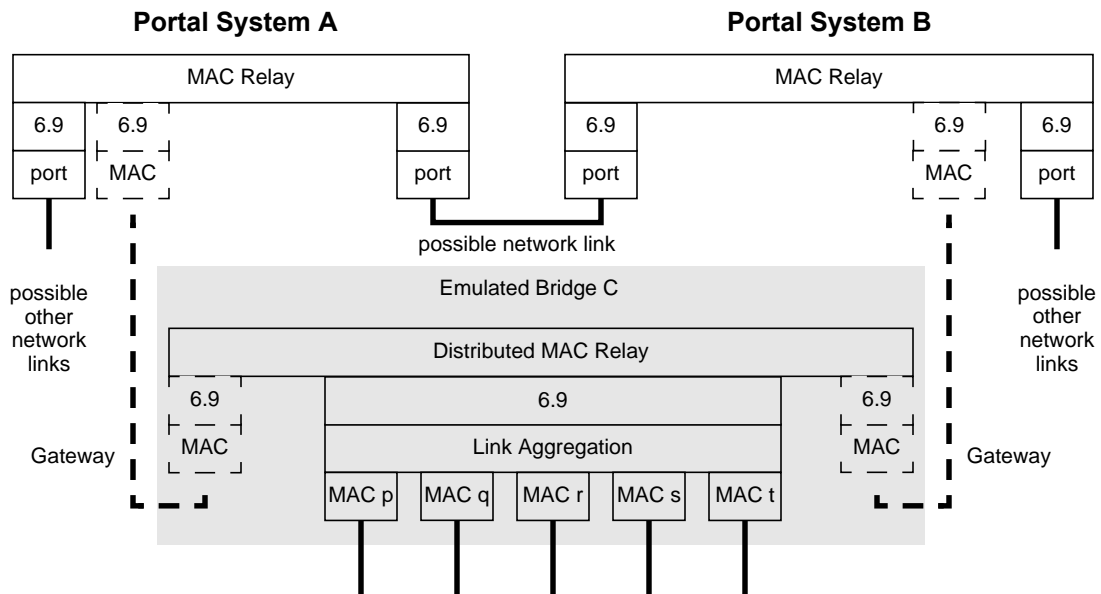


Figure E-1—Distributed VLAN Bridge as seen by other Bridges

Figure E-2 provides the Distributed MAC Relay: as seen by Bridge A and Bridge B.

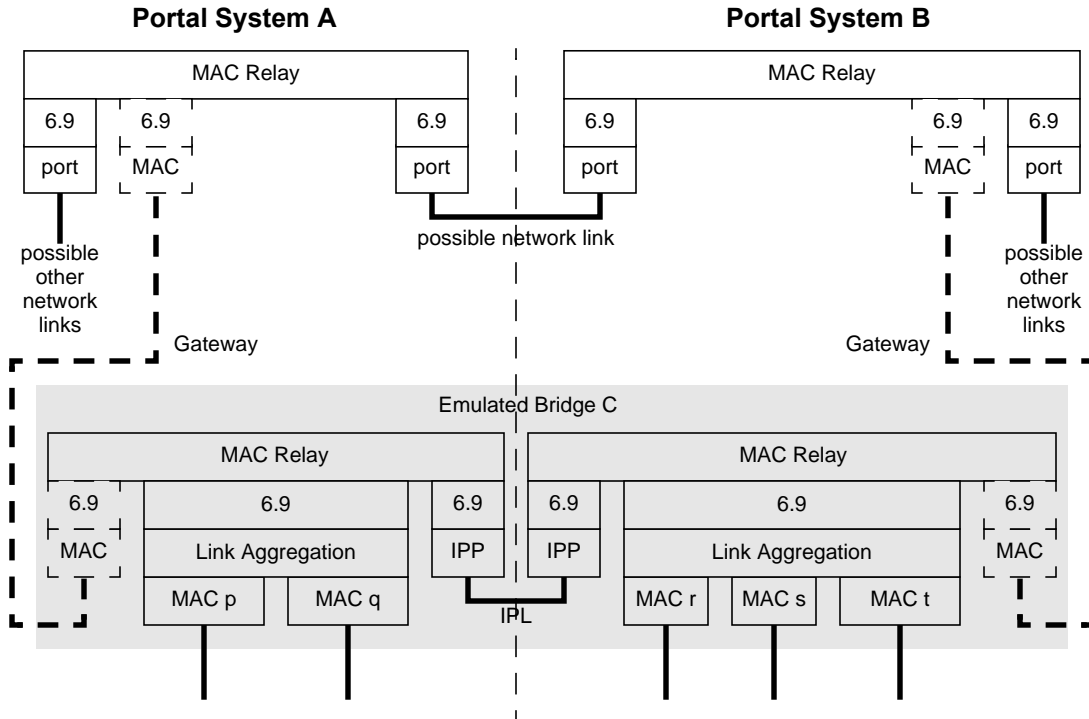


Figure E-2—Distributed MAC Relay: as seen by Bridge A and Bridge B

In this example the Gateway algorithm is based on VID, as required for VLAN Bridges implementing learning. In addition, we assume for simplicity that the Port algorithm is also based on VIDs. The forwarding rules that are specified in 9.3 can be satisfied in this case if the DR Function is implemented by the MAC Relay entity of Clause 8 in IEEE Std 802.1Q and every port on the MAC Relay Entity is supported by the EISS supporting function specified in Clause 6.9 of IEEE Std 802.1Q. The Gateway Conversation IDs on frames entering the DR Function through one of its ports are identified by the VID parameters on the EISS primitives that are invoked by the ISS primitives on the same DR Function's port. The Port Conversation IDs on frames on the Gateway or IPP are identified by the same VID parameters on the EISS primitives that are invoked by the ISS primitives on the Gateway or IPP respectively. The forwarding rules in 9.3 are implemented by configuring the VLAN memberships of each port on the DR Functions.

The example in Figure E-3 is similar to the previous one but instead of supporting the DRNI on a CNP or PNP, the DRNI is supported on a Customer Edge Port (CEP). Figure E-4 provides the Distributed PEB: as seen by Bridge A and Bridge B.

In this case the forwarding rules that are specified in 9.3 can be satisfied if the DR Function is implemented by the MAC Relay entity of Clause 8 in IEEE Std 802.1Q, the supporting functions on the IPPs and the Gateway to the MAC Relay Entity are the EISS supporting function specified in Clause 6.9 of IEEE Std 802.1Q, that is they are PNP ports, while the supporting functions of the Bridge Port for the Aggregator are those specified in Clause 15.4 of IEEE Std 802.1Q, that is the Aggregator is a CEP. The forwarding rules in 9.3 are implemented by configuring the S-VLAN memberships of each IPP and Gateway on the DR Function and the C-VID to S-VID mappings associated with the CEP configuration.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

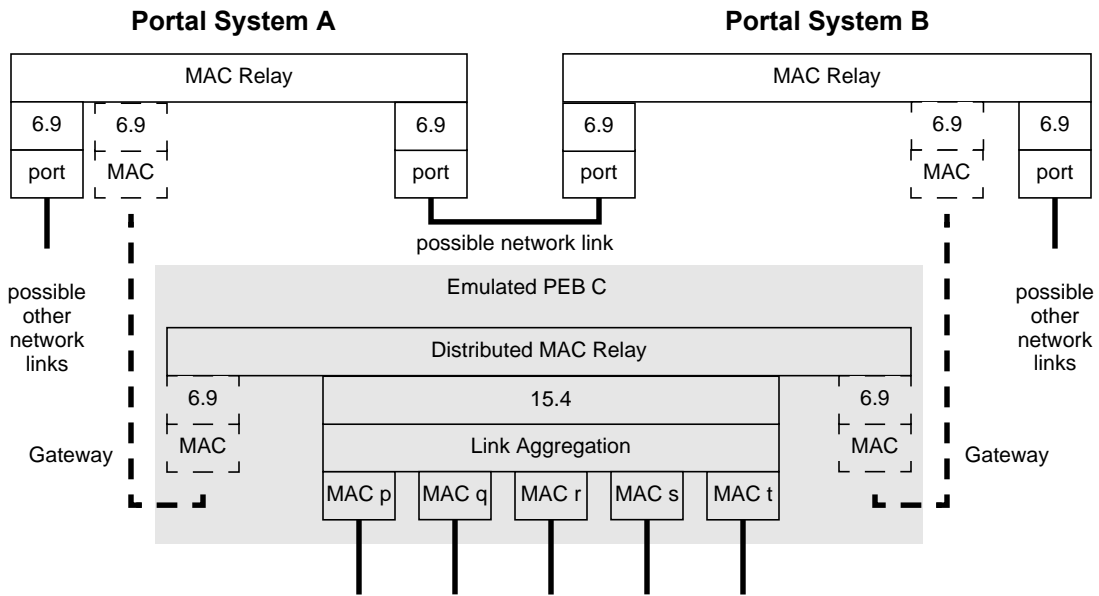


Figure E-3—DRNI on a Customer Edge Port

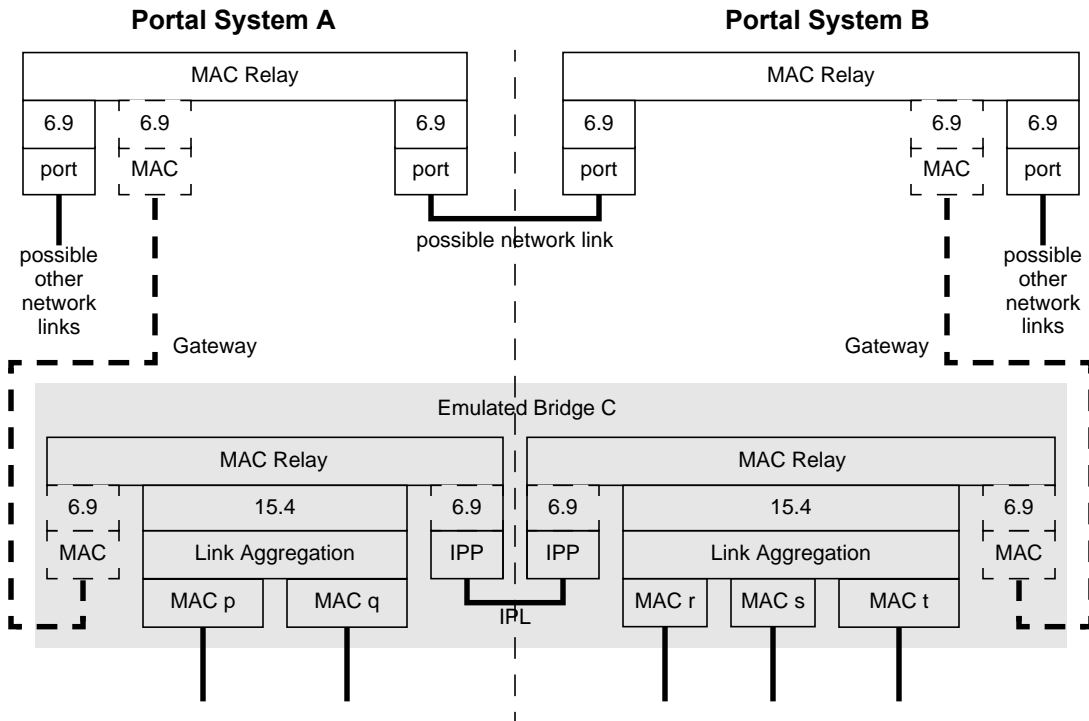


Figure E-4—Distributed PEB: as seen by Bridge A and Bridge B

Finally, the example in Figure E-5 is similar to the previous one but instead of supporting the DRNI on a CEP, the DRNI is supported on a Provider Instance Port (PIP) and the Port algorithm is based on I-SIDs.

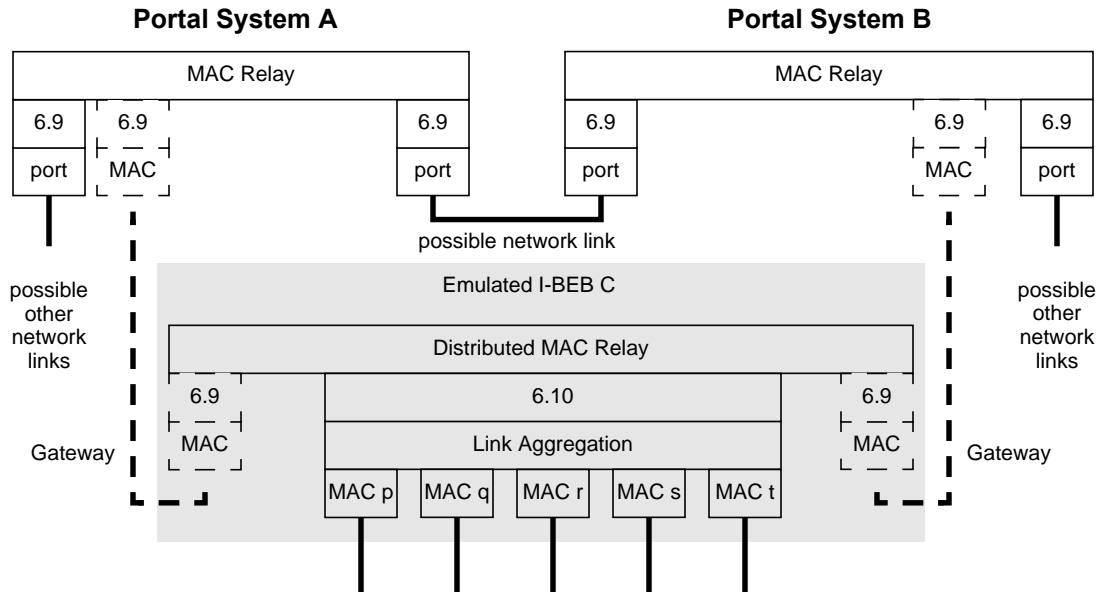


Figure E-5—DRNI on a Provider Instance Port

Figure E-6 provides the Distributed I-BEB: as seen by Bridge A and Bridge B.

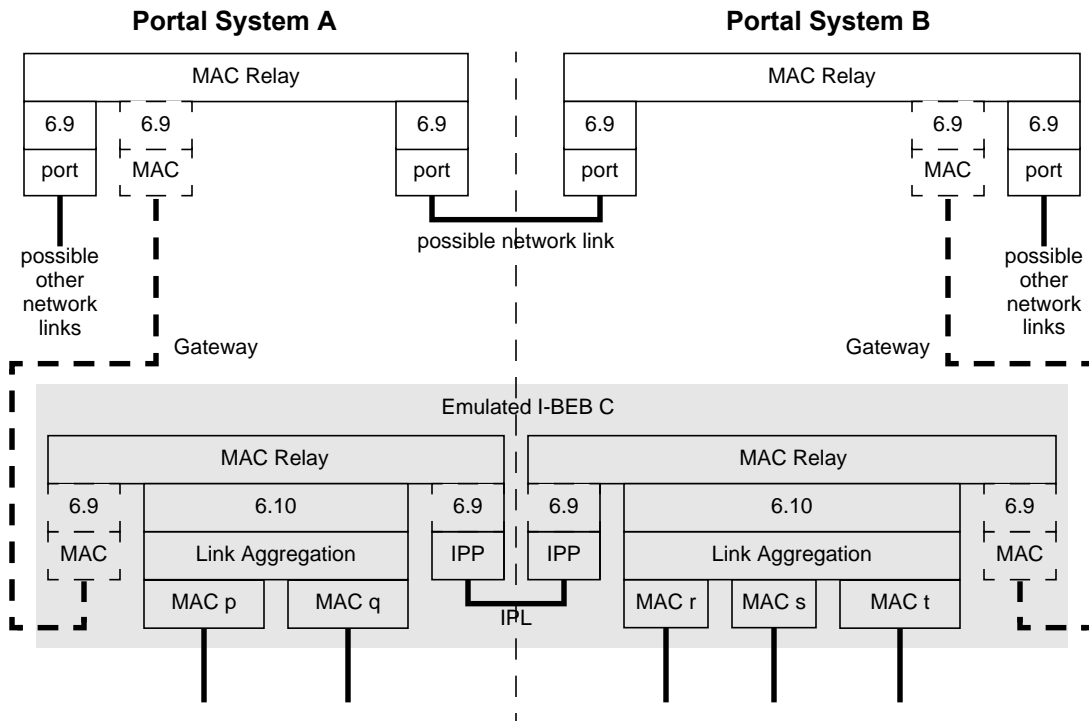


Figure E-6—Distributed I-BEB: as seen by Bridge A and Bridge B

1 In this case the forwarding rules that are specified in 9.3 can be satisfied if the DR Function is implemented
2 by the MAC Relay entity of Clause 8 in IEEE Std 802.1Q, the Gateway is implemented as a PNP, specified
3 in Clause 6.9 of IEEE Std 802.1Q, the Aggregator as a PIP, specified in Clause 6.10 of IEEE Std 802.1Q and
4 the IPPs as PNPs. The forwarding rules in 9.3 are implemented by configuring the S-VLAN memberships of
5 each IPP and Gateway on the DR Function and the S-VID to I-SID mappings and I-SID filtering associated
6 with the PIP configuration.
7

8 NOTE—When MAs are configured for monitoring services in the attached networks, the Portal System that hosts the
9 selected Gateway for a service is responsible for implementing an Up MEP monitoring that service (if such a MEP is
10 configured). How MEPs are implemented in the Distributed Relay, e.g. whether the functionality is distributed across
11 multiple physical ports or centralized somewhere in the Portal, is implementation dependent. The Intra-Portal Link is
12 outside the MA monitored by this MEP. The MEP functionality moves when the Gateway changes. The MEP state is not
13 synchronized during such a move.
14
15

16 **E.2 Higher Layer Entities in a Distributed Bridge**

17

18 Bridges support protocols [referred to as Higher Layer Entities in IEEE std 802.1Q (802.1Q-2011 subclause
19 8.2)] on each port of the bridge. Since the Emulated Bridge is seen by other bridges, it can be necessary or
20 desirable for the Emulated Bridge to participate in one or more of these protocols. It is theoretically possible
21 for the operation of such a protocol by an Emulated Bridge to be distributed among the Portal Systems, but
22 specifying that for each protocol is beyond the scope of this document. Recommended practice is for the
23 Portal System with the lowest Portal System Number to run these protocols on behalf of the Emulated
24 Bridge. This is referred to as the Host Portal System.
25

26 Higher Layer Entities typically communicate with neighboring bridges by transmitting and receiving frames
27 with a destination address that is one of the reserved addresses specified in IEEE std 802.1Q-2011 subclause
28 8.6.3. The protocols of the Emulated Bridge can transmit and receive such frames through any of its
29 Gateway ports and through its Aggregator port. Transmitting and receiving frames between the Emulated
30 Bridge and the Host Portal System is an operation completely internal to the Host Portal System.
31 Transmitting and receiving frames between the Emulated Bridge and the other Portal System(s) logically
32 occurs over the gateway, so the frames are carried on the IPL. Frames sent by the Emulated Bridge through
33 the Aggregator port can be transmitted on any Aggregation Port. It is generally easiest to select an
34 Aggregation Port that is a physical port of the Host Portal System. Frames received with a reserved
35 destination address on an Aggregation Port physically on the Host Portal System are forwarded to the
36 Emulated Bridge Higher Layer Entities as an internal operation. Frames received with a reserved destination
37 address on an Aggregation Port physically on a Portal System other than the Host Portal System are
38 forwarded on the IPL (unless they are intercepted by a parser multiplexer as described in 6.1.3).
39

40 Frames generated by Emulated Bridge protocols use the aDrniPortalAddr for the source address.
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Annex F

(normative)

Link Layer Discovery Protocol TLVs

F.1 Link Aggregation TLV

NOTE—The descriptions in this annex F.1 and its dependent subclauses supersede those in IEEE Std 802.1AB-2009, Annex E.8 and IEEE Std 802.1Q-2011, Annex D.2.7. All IEEE 802.1 Organizationally Specific TLVs conform to the LLDPDU bit and octet ordering conventions of 8.1 of IEEE Std 802.1AB. In addition, this annex clarifies and enhances the capabilities of LLDP with respect to Link Aggregation.

The Link Aggregation TLV indicates whether the link is capable of being aggregated, whether the link is currently in an aggregation, as specified in IEEE Std 802.1AX, whether the Aggregation Port from which the LLDPDU was transmitted is an Aggregation Port or an Aggregator, and if an Aggregation Port in an aggregation, the Port Identifier of the aggregation. Figure F-1 shows the format for this TLV.




Figure F-1—Link Aggregation TLV format

F.1.1 aggregation status

The link aggregation status field shall contain a bit map of the link aggregation capabilities and the current aggregation status of the link as defined in Table F-1.

Table F-1—Link aggregation capability/status

Bit	Function	Value/meaning
1	Aggregation capability	0 = not capable of being aggregated 1 = capable of being aggregated
2	Aggregation status	0 = not currently in aggregation 1 = currently in aggregation
3–4	Port type	43 bit number 00 = no port type specified 01 = transmitted from Aggregation Port 10 = transmitted from Aggregator 11 = transmitted from an Aggregator with a single Aggregation Port
5–8	reserved for future standardization	transmitted as 0, ignored on receipt

NOTE—The bits in Table F-1 are numbered 1–8. In the previous versions of this table in IEEE Std 802.1AB and IEEE Std 802.1Q, the bits were numbered 0–7. The numbering used here is compatible with the bit numbering used in the rest of 802.1AB, and with Figure F-1. This is an editorial correction only; there has been no change to the meanings of the bits as ordered by their proximity to the most-significant bit.

F.1.2 aggregated Port ID

The aggregated Port ID field shall contain the IEEE 802.3 aggregated Port Identifier, aAggPortID, derived from the ifNumber in the ifIndex for the interface.

NOTE—This field is meaningful only if the link concerned is currently a member of an Aggregation (Aggregation status = 1) and the LLDPDU was transmitted from an Aggregation Port (Port type = 00 or 01).

LLDP TLVs are sourced either from Aggregations Ports or Aggregators depending on the associated application specification.

F.1.3 Link Aggregation TLV usage rules

An LLDPDU should contain no more than one Link Aggregation TLV

In order to maximize both the utility of the Port type field, and interoperability with systems transmitting the Link Aggregation TLV without a Port type, any system transmitting a non-zero value for the Port type field of the Link Aggregation TLV on an Aggregator or any on of the Aggregation Ports associated with an Aggregator shall conform to the following rules. These rules apply only to LLDPDUs transmitted with the same destination MAC address.

- a) If a system transmits LLDPDUs from an Aggregator or one of its Aggregation Ports, and that LLDPDU has a non-zero Port type, then the system shall transmit a non-zero Port type in all of the LLDPDUs transmitted from all of the Aggregation Ports associated with that same Aggregator or from the Aggregator, itself, and each Port type shall indicate whether its LLDPDU was transmitted from the Aggregator, an Aggregation Port, or from a single-Port Aggregation.
- b) If an Aggregator is associated with a single Aggregation Port, it shall transmit LLDPDUs from either the Aggregator or from the Aggregation Port, but not both, and the Port type shall be 11 (Aggregator with single Aggregation Port).
- c) If an Aggregator is associated with more than one Aggregation Port, it shall transmit LLDPDUs with Port type 01 (Aggregation Port) only from Aggregation Ports, and with Port type 10 (Aggregator) only from the Aggregator.
- d) The LLDP instance associated with an Aggregation Port that is transmitting Port type 01 (Aggregation Port) shall not process LLDPDUs containing Port type 10 (Aggregator).
- e) The LLDP instance associated with an Aggregation Port that is transmitting Port type 10 (Aggregator) shall not process LLDPDUs containing Port type 01 (Aggregation Port).

This leaves undefined the question of what happens during the period (presumably brief) while an Aggregation is transitioning between having a single Aggregation Port and multiple Aggregation Ports.

F.1.4 Use of other TLVs on an Aggregator or Aggregation Link

An LLDPDU sent from an Aggregation Port can contain a different set of TLVs than an LLDPDU sent from higher layers through the Aggregator that includes that Aggregation Port. A given LLDPDU will be transmitted on one of the Aggregation Ports of the Aggregator. The LLDPDUs sent on an Aggregator are not

1 considered a conversation for the purposes of Link Aggregation; they may be distributed among the
2 Aggregation Ports arbitrarily.
3

4 In order to attach one or more instances of LLDP to an Aggregation Port, separately from attaching an
5 instance of LLDP as a client of the Link Aggregation sublayer, an LLDP Parser/Multiplexer shall be used.
6

7 Any system transmitting a non-zero value for the Port type field of the Link Aggregation TLV shall use an
8 LLDP Parser/Multiplexer (6.1.3) to attach the instance(s) of LLDP to the Aggregation Port, and shall
9 conform to the following rules for other TLVs defined in IEEE Std 802.1AB:
10

- 11 a) The Port ID TLV (IEEE Std 802.1AB-2009 clause 8.5.3) transmitted in an LLDPDU with a Link
12 Aggregation TLV Port type 10 (Aggregator) shall have a different port ID than the port IDs in any
13 LLDPDUs transmitted from its component Aggregation Ports carrying a Link Aggregation TLV
14 Port type 01 (Aggregation Port).
 - 15 b) The following TLVs shall not be transmitted in an LLDPDU carrying a Port type (from Table F-1) of
16 10 (Aggregator):
 - 17 1) MAC/PHY Configuration Status (802.1AB-2009 Clause F.2).
 - 18 2) Power Via Medium Dependent Interface (MDI) (802.1AB-2009 Clause F.3).
- 19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Annex G

(normative)

Network / IPL sharing by time - MAC Address synchronization

The circumstances under which MAC address synchronization is required are illustrated in Figure G-1 and summarized below:.




Figure G-1—MAC address synchronization

- a) The method that is used to distinguish frames on a network link from frames on a particular IPL is the Network / IPL sharing by time (9.3.2.1);
- b) The Portal Systems are Bridges;
- c) There is no agreement on symmetric Port Conversation IDs across the DRNI;
- d) The IPL (B to Y) is unblocked in the network topology, so cannot serve as a dedicated IPL.

Then, when Bridge B receives a frame with an unknown source address from a port other than the IPL (e.g. from A, as shown, or from the Aggregator Port), it becomes the Gateway for that frame, and emits the frame over the IPL to enable the normal source learning process. However, when the frame reaches the Bridge at the other end of the IPL, the location of the source is now ambiguous;

- e) It could have come via a network port on Bridge B (as shown), in which case the correct response by Bridge Y to receiving that learned address as the destination of a frame is to forward it over the IPL
- f) It could however have been received at Bridge B over the Aggregator Port, in which case the correct response by Bridge Y to receiving, on a non-IPL network port, the learned address as the destination of a frame is to consider itself to be the Gateway, and forward the frame over the Aggregator Port.

In the absence of address synchronization, the only way to ensure connectivity is for a Bridge to replicate frames arriving on a non-IPL network port and carrying a destination address which has been learned via the IPL, and forward one copy over the IPL, and one copy over the Aggregator Port.

A Bridge shall perform the following actions to ensure that frame duplication does not occur whilst in this frame replication mode:

- g) A frame received over the IPL shall never be forwarded over the Aggregator Port.
- h) A frame received over the IPL with a destination address which was learned from the Aggregator Port shall be discarded.

1 NOTE—These actions ensure that frame duplication cannot occur, and limit the propagation of replicated frames as far
2 as is practicable in the absence of address synchronization. DRNI operation is correct, but will exhibit degradation due to
3 unnecessary propagation of individually-addressed frames.
4
5

6 **G.1 Address Synchronization - design goals**

7

8 The address synchronization process has the following goals:
9

- 10 a) It should be simple and be capable of high performance implementation, to allow bursts of learning
11 when stations are added to the network.
- 12 b) It should operate as far as possible as an adjunct to the normal MAC learning process.
13

14 **G.2 Address Synchronization - non-goals.**

15

- 16 a) The process does not need to be totally reliable, because an infrequent delays in synchronization will
17 cause only temporarily degraded operation of the DRNI, not aberrant operation.
18
19

20 **G.3 Protocol Summary**

21

22 A Bridge receiving a frame from non-IPL ports with an unknown source address performs the MAC
23 learning procedure in the normal way (Clause 8.7 in IEEE Std 802.1Q-2011). The Bridge shall also form or
24 extend an Address Sync TLV binding the learned MAC to its port of arrival, either the Aggregator Port or a
25 non-IPL network port. Periodically, one or more Address Sync TLVs are packed into an ASPDU and
26 forwarded on the IPL only.
27

28 A Bridge receiving frames from non-IPL ports shall also respond to the receipt from the IPL of an ASPDU
29 containing one or more Address Request TLVs, by replying with an Address Sync TLV for the MACs
30 requested in each Address Request TLV and reporting the port of arrival (Aggregator Port or a non-IPL
31 network port) of each MAC.
32

33 A Bridge receiving an ASPDU from the IPL parses each Address Sync TLV. If there is no Dynamic
34 Filtering Entry (Clause 8.8.3 in IEEE Std 802.1Q-2011) corresponding to a MAC address in the TLV, no
35 further processing is performed for that MAC entry. If there is such an entry, the Dynamic Filtering Entry is
36 annotated with the reported source of the frame at the other Bridge, so that the receiving Bridge can
37 correctly direct frames destined to the learned address, either to its Aggregator port or the IPL as
38 appropriate.
39

40 If a Bridge has created a Dynamic Filtering Entry (Clause 8.8.3 in IEEE Std 802.1Q-2011) for a MAC
41 learned from the IPL, but no Address Sync TLV has been received for that MAC from the IPL within a
42 defined period, the Bridge may issue an ASPDU towards the IPL containing one or more Address Request
43 TLVs soliciting source port information for such learned MACs.
44

45 The reported source port attributes associated with learned MAC addresses are never explicitly aged. The
46 normal ageing process for Dynamic Filtering Entries (Clause 8.8.3 in IEEE Std 802.1Q-2011) removes both
47 MAC address and any reported source port attributes associated with it.
48
49

50 **G.4 Address Synchronization Description**

51

52 The protocol is tightly coupled to the normal MAC learning procedures. A Bridge receiving a frame with an
53 unknown source MAC address from a non-IPL port (including the Aggregator Port):
54

- a) Initiates the normal source learning process (Clause 8.8.3 in IEEE Std 802.1Q-2011).
- b) For each MAC address newly learned on the Aggregator port, it includes it in an Address Sync TLV, recording the VLAN on which the address was received, the value of the source MAC address, and an indication of its source as the Aggregator Port.
- c) Optionally, for each MAC address newly learned on a non-IPL network port, it includes it in an Address Sync TLV, recording the VLAN on which the address was received, the value of the source MAC, and an indication of its source as a non-IPL network port.

NOTE 1—Optionality allows control of the ASPDU overhead in environments with significant learning loads; the heavier learning load is assumed to be on-net.

- d) It concatenates multiple such TLVs created at b) and c) above, if available, into a ASPDU.
- e) When the ASPDU is full or after the time *ASPDU_Transmit_Holdoff* (default 100ms) from the assembly of the earliest received MAC entry in the ASPDU (whichever is sooner), the ASPDU is forwarded onto the IPL, irrespective of the number of new TLVs it then contains, subject to the constraint that *ASPDU_Transmit_Min_Period* (default 10ms) is not infringed. Under conditions of very high transient learning load, ASPDUs may back up. If this occurs, the unsent ASPDUs should be buffered in a FIFO queue and sent, eldest first, as rapidly as allowed.

NOTE 2—At 1Gbps, 1440 Bytes requires ~12μs to transmit, so ASPDUs are limited to ~0.1% link capacity; More could probably be tolerated, but overrunning the receiver processing should be avoided.

- f) Transmission of ASPDUs containing newly learned MAC addresses may be repeated *ASPDU_Repeat_Number* (default 0) times after successive intervals of *ASPDU_Fast_Repeat* (default 4 x *ASPDU_Transmit_Holdoff*), to allow rapid recovery from a lost ASPDU.

A Bridge receiving frames from the IPL receives and parses ASPDUs. Address Request TLVs are processed as described in items g) and h) below:

- g) For each MAC in the Address Request TLV, a corresponding MAC entry in an Address Sync TLV is formed, including the source port (Aggregator or a non-IPL network port) for that MAC, or “Aged” if no Dynamic Filtering Entry (8.8.3) exists for that MAC.
- h) Address Sync TLVs are concatenated into an ASPDU and forwarded towards the IPL.

Address Sync TLVs are processed as described in item i) below:

- i) For each MAC entry carried in an Address Sync TLV received:
 - 1) if there is no Dynamic Filtering Entry (8.8.3 in IEEE Std 802.1Q-2011) corresponding to that MAC learned over the IPL, the MAC entry is discarded; (this covers the circumstance of “not yet learned” AND “Aged”; in particular, the recipient ages learned MAC addresses according to the normal reverse path learning rules)
 - 2) else the source port data in the MAC entry is used to overwrite the source attribute for that MAC.
- j) If a Dynamic Filtering Entry corresponding to a MAC address learned through the IPL has not received a corresponding Address Sync TLV after a time since its initial creation of 1.5 x *ASPDU_Transmit_Holdoff*, the Bridge may form an Address Request TLV related to that MAC address in order to solicit that information from its Neighbor Bridge.
- k) Address Request TLVs formed by the process j) above may be included in ASPDUs, together with any available Address Sync TLVs, and be forwarded towards the IPL. No second or subsequent Address Request for any MAC address shall be made until time *ASPDU_Fast_Repeat* has elapsed since the previous request.
- l) If a Dynamic Filtering Entry corresponding to a MAC address learned through the IPL has not received a MAC entry in an Address Sync TLV relating to that entry within *Default_Network_Port_Timeout* (default = 3 x *ASPDU_Transmit_Holdoff* +

1 *ASPDU_Repeat_Number* x *ASPDU_Fast_Repeat*) of its initial creation, the Bridge may set the
2 source attribute for that MAC entry to be the non-IPL network port of its Neighbor Bridge.
3

4 NOTE 3—Requires only transmission of ASPDUs for MACs learned on the Aggregator port, at the expense of longer
5 learning times for non-IPL network port associations.
6

7 **G.5 ASPDU transmission, addressing, and protocol identification**

8
9
10 Address Synchronization Protocol Data Units (ASPDUs) are transmitted and received using the service
11 provided by an LLC entity that uses, in turn, a single instance of the MAC Service provided at an MSAP
12 associated with an IPP. Each ASPDU is transmitted as a single MAC service request, and received as a
13 single MAC service indication, with the following parameters:
14

- 15 a) destination address (G.5.1)
- 16 b) source address (G.5.2)
- 17 c) MSDU
- 18 d) priority (G.5.3)
- 19

20 The MSDU of each request and indication comprises a number of octets that provide EtherType protocol
21 identification (G.5.4) followed by the ASPDU proper (G.5.5).
22

23 NOTE—The complete format of an address synchronization frame depends not only on the ASPDU format, as specified
24 in this clause, but also on the media access method dependent procedures used to support the MAC Service.
25

26 **G.5.1 Destination MAC Address**

27
28
29 The destination address for each MAC service request used to transmit a ASPDU shall be the same as the
30 destination MAC address used by the DRCP protocol which is selected by IPP Managed Objects Class
31 (7.4.2). Its default values shall be the Nearest non-TPMR Bridge group address.
32
33

34 **G.5.2 Source MAC Address**

35
36 The source address for each MAC service request used to transmit an ASPDU shall be an individual address
37 associated with the IPP MSAP at which the request is made.
38
39

40 **G.5.3 Priority**

41
42 The priority associated with each MAC Service request should be the default associated with the IPP MSAP.
43
44

45 **G.5.4 Encapsulation of ASPDUs in frames**

46
47 An ASPDU is encoded in the *mac_service_data_unit* parameter of an *M_UNITDATA.request* or
48 *M_UNITDATA.indication*. The first octets of the *mac_service_data_unit* are a protocol identifier, followed
49 by the ASPDU, followed by padding octets, if any, as required by the underlying MAC service.
50

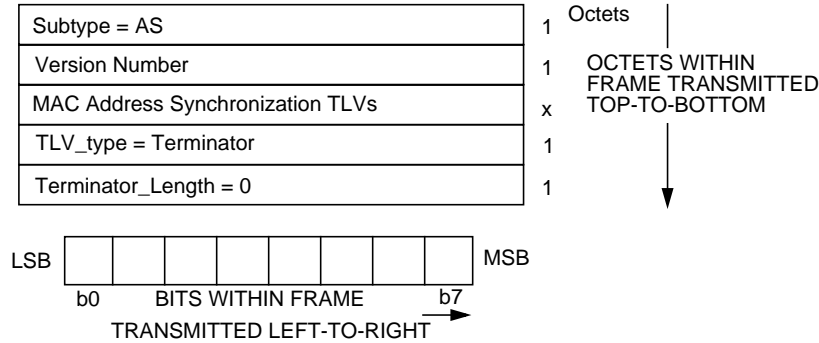
51 Where the ISS instance used to transmit and receive frames is provided by a media access control method
52 that can support EtherType encoding directly (e.g., is an IEEE 802.3 MAC), the protocol identifier is two
53
54

1 octets in length. All ASPDUs are identified by the DRNI EtherType specified in Table 9-4, which is
2 followed immediately by the structure of Figure G-2.
3
4

5 G.5.5 ASPDU structure and encoding

6 G.5.5.1 ASPDU structure

7 All ASPDUs comprise an integral number of octets and are represented following the rules described in
8 9.4.3.1. The ASPDU structure shall be as shown in Figure G-2 and as further described in the following field
9 definitions:
10
11



27 **Figure G-2—ASPDU structure**

- 28
29
30
31
32
33
34
35
36
37
38
39
- a) *Subtype*. The Subtype field identifies the specific Protocol being encapsulated. ASPDUs carry the Subtype value 0x02.
 - b) *Version Number*. This identifies the MAC Address Synchronization protocol version; implementations conformant to this version of the standard carry the value 0x01.
 - c) *MAC Address Synchronization TLVs*. There are two MAC Address Synchronization TLVs, which may be mixed freely into an ASPDU for transmission as required:
 - 1) Address Sync TLV (G.5.5.1.1);
 - 2) Address Request TLV (G.5.5.1.2).
 - d) *TLV_type = Terminator*. This field indicates the nature of the information carried in this TLV-tuple. Terminator (end of message) information is identified by the integer value 0x00.
 - e) *Terminator_Length*. This field indicates the length (in octets) of this TLV-tuple. Terminator information uses a length value of 0 (0x00).

40 NOTE—The ASPDU TLV Length fields provide the length of the TLV Value fields in contrast to the LACP TLV Length
41 fields which provide the length of the total TLV (including the 2 octets for the TLV Type and TLV Length field). This
42 enables the use of TLVs carrying a TLV Value field of a length up to 256 octets.
43
44
45
46
47
48
49
50
51
52
53
54

G.5.5.1.1 Address Sync TLV

The Address Sync TLV structure shall be as shown in Figure G-3 and as further described in the following field definitions:




Figure G-3—Address Sync TLV

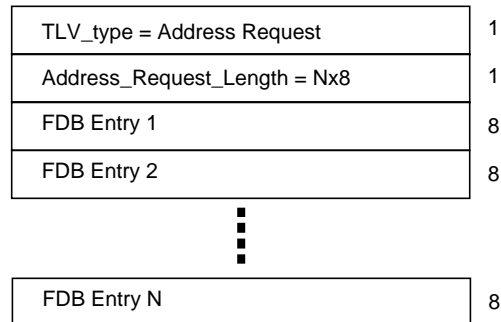
- a) *TLV_type = Address Sync*. This field indicates the nature of the information carried in this TLV-tuple. The Address Sync TLV is identified by the integer value 0x01.
- b) *Address_Sync_Length*. This field indicates the length (in octets) of this entire structure, including the header, of this TLV-tuple. The Address Sync TLV uses a length value of Nx8 where N can be 1 to a maximum of 31.
- c) *FDB Entry*. This field contains the FDB entries carried by an ASPDU. Its structure is provided by Table G-1 below.

Table G-1—FDB Entry

Parameter	Size
Reserved	2 bits
FDB Entry Aged = 0 / FDB Entry Valid =1	1 bit
Non-IPL network port (=0) / Aggregator (=1)	1 bit
VID	12 bits
MAC Address	6 Octets
	8 Octets

G.5.5.1.2 Address Request TLV

The Address Request TLV structure shall be as shown in Figure G-3 and as further described in the following field definitions:

**Figure G-4—Address Request TLV**

- a) *TLV_type = Address Request*. This field indicates the nature of the information carried in this TLV-tuple. The Address Request TLV is identified by the integer value 0x02.
- b) *Address_Request_Length*. This field indicates the length (in octets) of this entire structure, including the header, of this TLV-tuple. The Address Request TLV uses a length value of Nx8 where N can be 1 to a maximum of 31.
- c) *FDB Entry*. This field contains the FDB entries carried by an ASPDU. Its structure is provided by Table G-1 above. In the Address Request TLV, each FDB Entry Valid field and each Source Network Port field are set to zero on transmission and ignored on receipt.