

1 In reviewing Clause 6.6 of P802.1AX-Rev/D0.3 I found significant problems with this mainly old text,
2 possibly as a result of its incremental development. Coming to it fresh I found that the high level view that is
3 hinted at in the initial clauses is mainly missing, that the division of material into these initial clauses
4 resulted in much repetition, that some of the statements made are more wishful than true, and I couldn't
5 figure out how it is all meant to work. I can't see the wood for the trees. The degree of reshuffling that I think
6 is necessary makes it very difficult to write a set of point changes that an editor would have a chance of
7 reassembling into a coherent whole, and in any case I needed to write the intended result to debug the
8 proposed changes.

9 The result follows. It is intended as a complete replacement for the 6.6 introductory text, and the following
10 clauses 6.6.1, 6.6.1.1, 6.6.2 and for 6.6.3 up to pg 85 line 17 "The parameters ..". I am uncertain what should
11 be said after this point. I have the feeling that some of the detail is chasing unrealistic goals, but can't be
12 sure.

13

1 6.6 Conversation-sensitive Collection and Distribution

2 Conversation-sensitive Distribution and Collection (CSCD) allows administrative control of the Frame
3 Distributor's selection of the Aggregation Link for each frame, and allows the Collector to accept only
4 frames received on the expected Aggregation Link.

5 The distribution of frames can be controlled whether or not the Link Aggregation partner supports CSCD. If
6 the partner has a matching CSCD configuration, each conversation supported by the LAG can be distributed
7 to the same Aggregation Link in both directions. This facilitates load balancing, detailed traffic management
8 and bandwidth control, and implementation of monitoring and policing functions on a LAG.

9 Conversation-sensitive Distribution and Collection does not require a particular distribution algorithm;
10 rather it creates a structure for identifying and configuring a variety of standardized or vendor-defined
11 distribution algorithms. It uses the following:

- 12 a) A Port Algorithm (identified by the Actor Port Algorithm variable) used to associate each frame
13 transmitted or potentially received by an Aggregator with a 12-bit Port Conversation ID (6.6.1).
- 14 b) A Distribution Conversation Mask for each Aggregation Port, with a single bit for each Port
15 Conversation ID that permits or denies distribution to the port's Aggregation Link.
- 16 c) A Collection Conversation Mask for each Aggregation Port, with a single bit for each Port
17 Conversation ID that permits or denies reception from the port's Aggregation Link.

18 The Distribution and Collection Conversation Masks for each Aggregation Port are automatically updated as
19 a consequence of the configuration of, or changes to, the following:

- 20 d) LACP Selected (6.4.7), Actor_Oper_State.Synchronization (6.4.2.3), and Wtr_Waiting (6.4.6).
21 Only Aggregation Ports that are correctly attached to the same Aggregator and are not in the Mux
22 machine ATTACHED_WTR state (6.4.14) influence that Aggregator's collection and distribution.
- 23 e) A Link Number (6.6.2) for each such 'active' [d) above] Aggregation Port.
- 24 f) An Admin Conversation Link Map for the attached Aggregator, that determines which (if any) of
25 the 'active' Aggregation Ports (identified by their Link Numbers) distribute frames for any given
26 Port Conversation ID .
- 27 g) LACP Actor_Oper_State.Distributing and Actor_Oper_State.Collecting for the Aggregation Port.
- 28 h) A Discard Wrong Conversation (DWC,<ref>⁶) variable for the attached Aggregator.

29 The Admin Conversation Link Map and the active Link Numbers [e) and f) above] are used to calculate a
30 Port Conversation Vector for the Aggregator, mapping each Port Conversation ID to an Aggregation Port
31 Number. The Port Conversation Vector is then used to generate a Port Conversation Mask for each of the
32 Aggregator's active Aggregation Ports, with a single bit for each Port Conversation ID. The bit is set
33 (TRUE) if frames associated with the Port Conversation ID are to be distributed to that Aggregation.

34 If an Aggregation Port's Actor_Oper_State.Distributing is FALSE, its Distribution Conversation Mask is
35 cleared, inhibiting transmission. Otherwise it is updated using its Port Conversation Mask, taking care not to
36 permit transmission of any conversation through more than one Aggregation Port at the same time.

37 CSCD-capable Systems use Version 2 LACP to exchange the Actor's and Partner's Port Algorithm
38 identifiers and a digest of their Admin Conversation Link Map's, and to agree a common Link Number for
39 each Aggregation Port, so every Actor can know if its Partner is (or is about to) distribute the same
40 conversations to the same links. If they do not differ and the Aggregator's DWC variable value is AUTO,
41 each Aggregation Port's Port Conversation Mask is used to update its Collection Conversation Mask.

⁶Reference required, stating the three possible values and possibly more detail. To the management section description?

1 If the Actor and Partner's distribution configuration differs or the DWC value is FORCE_FALSE, the
 2 Collection Conversation Masks are updated to permit reception of any conversation on any of the
 3 Aggregation Ports. The DWC value can also be FORCE_TRUE, with Collection Conversation Mask
 4 updating independent of Port Algorithm or Conversation Link Map differences.

5 **6.6.1 Port Algorithms and Port Conversation IDs**

6 A Port Algorithm specifies how each frame (i.e. service requests from the Aggregator Port and service
 7 indications from the Aggregation Port) is associated with a Port Conversation ID taking a value between 0
 8 and 4095 and is identified by the Actor Port Algorithm variable.

9 The three most significant octets of the Actor Port Algorithm are an OUI identifying the organization
 10 responsible for the specification of the algorithm. The fourth octet differentiates algorithms specified by the
 11 same organization. The specification defines the frame field(s) used, and how their contents are mapped to
 12 specific Port Conversation ID values. Clause 8.1 specifies a number of algorithms.

13 Possible algorithms map a frame field (e.g. a VLAN identifier) directly to the Port Conversation ID, or use
 14 an algorithmic mapping (e.g. a hash algorithm that reduces the contents of specified fields to a 12-bit
 15 integer). This standard also supports table-driven mappings (e.g. of an I-SID to a Port Conversation ID) and
 16 algorithmic and table-driven mapping combinations, by specifying an optional per Aggregator configurable
 17 Aggregator Admin_Conversation_Service_ID_Map that maps 32-bit Service IDs to Port Conversation IDs.
 18 A CSCD-capable LACP Actor and its Partner can exchange an MD5 digest of their
 19 Admin_Conversation_Service_ID_Maps as well as their Actor Port Algorithm identifiers to confirm that
 20 they are using the same frame to Port Conversation ID mapping.

21 **6.6.2 Link Numbers and Link Selection**

22 A CSCD-capable System's Admin Conversation Link Maps (one per Aggregator) use Link Numbers to
 23 identify Aggregation Links, with a local mapping of each Map's Link Numbers to the Actor_Port_Numbers
 24 used to identify Aggregation Ports within the System as a whole. The Version 2 LACPDUs transmitted by
 25 each Aggregation Port include a locally configured Admin_Link_Number that is unique among the
 26 Aggregation Ports with the same Aggregation Key (because they could be attached to the same Aggregator).

27 If all the following conditions apply:

- 28 a) an Aggregation Port is in current receipt (i.e. Actor_Oper_Port_State.Defaulted and
 29 Actor_Oper_Port_State.Expired are both FALSE) of a Version 2 LACPDU
- 30 b) the attached link is active as described in 6.6 d) and e) above (i.e. Selected and
 31 Actor_Oper_State.Synchronization are TRUE and Wtr_Waiting is FALSE)
- 32 c) the Actor has a lower priority than its Partner (<ref>⁷)
- 33 d) the Actor and its Partner are using the same Port Algorithm (with the same
 34 Admin_Conversation_Service_Map if used by the algorithm) and the same Admin Conversation
 35 Link Map (<ref>⁸)

36 then the operational Link_Number used by the Aggregation Port in conjunction with its Admin
 37 Conversation Link Map is the Admin_Link_Number transmitted by the Actor's Partner
 38 (Partner_Link_Number). Otherwise the Actor's Admin_Link_Number is used.

⁷The reference to specify that the concatenation of the Partner_Oper_System_Priority, Partner_Oper_System,
 Partner_Oper_Port_Priority, and Partner_Oper_Port_Number is numerically less than the concatenation of
 Actor_Port_System_Priority, Actor_Port_System, Actor_Port_Priority, and Actor_Port_Number.

⁸The reference to specify that the values of Actor_Port_Algorithm and Partner_Port_Algorithm, of
 Actor_Conversation_Service_Mapping_Digest and Partner_Conversation_Service_Mapping_Digest, and of
 Actor_Conversation_Link_Digest and Partner_Conversation_Link_Digest match respectively, possibly by using the Differ variables.

1 Each Link_Map is configured with a link selection priority list, a sequence of one or more Link_Numbers in
2 preference order (highest to lowest), for each Port Conversation ID. The Link_Number of the first currently
3 active link in each list identifies the Aggregation Port to which frames associated with that Port
4 Conversation ID are distributed.