CommentNumber: 1
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 3/43
Subclause: 3.5/43.4.2.2/43.5.3.2
Page: ??/119/142
Line: ??/21/50
CommentType (E, T, ER, or TR): TR
Comment:

43.4.2.2 Page 119 Line 21 and 43.5.3.2 Page 142 Line 50 correctly state that LACP and Marker PDUs shall not be Tagged, and refers the reader to Clause 3, which (since publication of .3ac) now talks about the Q Tagged Frame format. This is an important requirement, in order to preserve the integrity of the 802.3 and VLAN Bridging architectures, and applies not only to Link Aggregation exchanges but also MAC Control & potentially other internal 802.3 protocols that are yet to be invented. I believe that there should be a blanket statement in Clause 3 that covers LACPDUs, Marker PDUs, MAC Control frames, and any other frames that we may in future invent for the purposes of the internal operation of the 802.3 MAC, stating that such frames shall not be Tagged, so that there is no possibility of confusion or misinterpretation of this requirement now or in the future. It is easier to fix this potential problem now by making Clause 3 crystal clear on this point than it will be if misinterpretation of Clause 3 leads one or more implementors to find creative uses for Tagged versions of these frames.

CommentEnd:

SuggestedRemedy:

Add the following paragraph at the start of 3.5 Elements of the Tagged MAC frame, before the existing first para:

"Tagged MAC frames are frames that carry a QTag Prefix; these are frames generated by the MAC Client in accordance with the requirements stated in IEEE Std 802.1Q-1998. MAC frames that are sourced by, and are destined for, protocol entities that exist within the 802.3 MAC sublayer (for example, MAC Control frames, LACPDUs and Marker frames), all make use of the basic IEEE 802.3 frame format; such frames shall not be Tagged."

Leave the existing statements in 43.4.2.2 and 43.5.3.2 as they are; strictly speaking, the Clause 3 statement is sufficient, but re-iterating the requirement reinforces the point.

Consider retro-fitting a similar statement into the definition of MAC Control.

RemedyEnd:
Editor'sRecommendationStart:

Discuss. The particular solution offered doesn't precisely work, since MAC Control frames, LACPDUs, and Marker frames are NOT sourced or destined for entities within the MAC sublayer (all belong to clients of the 802.3 MAC).

The fundamental issue is considerably larger, i.e., what is the architecture of an end station with regard to VLANs and tagging. This is not directly addressed by 802.1Q, being a standard for VLAN-aware bridges (not end stations). I agree that LACPDUs, Marker frames, and PAUSE frames (the only currently-defined MAC Control frames) should not be tagged, but it is not clear that there will never be a MAC Control frame for which tagging could be useful.

A blanket change to Clause 3 is out of the scope of this Task Force. I also believe that it is not a "trivial" change that should be done as a "service to humanity".

Editor'sRecommendationEnd:
ResolutionStart:

Withdrawn pending.

ResolutionEnd:

CommentNumber: 2
CommenterName: Graham Short
CommenterEmail: Graham.Short@tiuk.ti.com
CommenterPhone: +44-1604-663410
CommenterCompany: Texas Instruments
Clause: 30
Subclause: 30.7.1.7 and .8
Page: 19
Line: 2 and 16
CommentType (E, T, ER, or TR): E
Comment:
There are references to section 43.2.11.4 here. Not sure what the reference is supposed to refer to but this section does not exist.
CommentEnd:

SuggestedRemedy:
Correct reference
RemedyEnd:
Editor'sRecommendationStart:

Change reference to 43.6.2 in both cases.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 3
CommenterName: Keith.Balmer

CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 30
Subclause: 30.7.1.14
Page: 20
Line: 54
CommentType (E, T, ER, or TR): E
Comment:
"...collecting and/or distributing,...".
Cannot be 'or' (i.e. only distributing).
CommentEnd:

SuggestedRemedy:
Remove "and/or distributing".
Ditto p21 line 1.
RemedyEnd:
Editor'sRecommendationStart:

Change to "collecting, or both collecting and distributing"

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 4
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 30
Subclause: 30.7.1.25
Page: 24
Line: 23
CommentType (E, T, ER, or TR): E
Comment:
Typo. Replace "fram" with "frame".
CommentEnd:

SuggestedRemedy:
Replace "fram" with "frame".
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:

ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 5
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 30
Subclause: 30.7.1.31
Page: 26
Line: 1
CommentType (E, T, ER, or TR): E
Comment:
Syntax described as "Same as aAggID".
aAggID says "INTEGER". Why not just say "INTEGER" here?
CommentEnd:

SuggestedRemedy:
Change "Same as aAggID" to "INTEGER".

Ditto for 30.7.1.32 (page 26, line 16)
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 6
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 30C
Subclause: All
Page: 67 on
Line: All
CommentType (E, T, ER, or TR): TR (Procedural)
Comment:

Having taken the radical (for 802.3) step of defining an SNMP MIB, we

should put in place the procedural steps needed to ensure that what we produce ends up as the accepted MIB both in IEEE and the IETF. In particular, we must make sure that we don’t end up in the position where the IETF feels the need to "tweak" what we define in order for it to be acceptable to them, thereby generating two different "standard" MIBs. The right approach with this seems to be to pre-handle any such problems by making sure that we develop a MIB that satisfies both our requirements and any requirements that may be expressed by interested parties in the IETF.

CommentEnd:

SuggestedRemedy:

In a similar vein to the current proposals regarding Internationalisation of 802 standards, 802.3 should invite the relevant IETF working group membership to comment on the 802.3ad MIB definition at both the WG ballot and the Sponsor Ballot stages, and their comments should be taken into consideration in the usual way as part of the process of ballot resolution; they should also be invited to take part in the editing meetings if they feel so moved. This will ensure that any IETF concerns are handled in a timely manner. Following Sponsor Ballot approval, we take the necessary steps to generate an informational RFC in the IETF; this would point at the IEEE standard for the MIB definition, and describe the process of its development (in particular, it would identify the fact that the MIB was the output of collaborative work between 802.3 and IETF working groups). A Note or a footnote in the .3ad MIB text should also identify the collaborative development process.

In addition to commenting on the MIB definition itself, the relevant IETF WG should be invited to propose the appropriate place in the MIB tree for the .3ad MIB to be attached (the choice would be between a specific 802.3 arc or a MIB 2 arc).

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Add appropriate IETF designees to WG ballot pool for non-voter comment submissions. Work with IEEE to see if the MIB portion of this standard can be made available on-line/free in lieu of producing an IETF RFC.

ResolutionEnd:

CommentNumber: 7
CommenterName: Graham Short
CommenterEmail: Graham.Short@tiuk.ti.com
CommenterPhone: +44-1604-663410

CommenterCompany: Texas Instruments
Clause: 30C
Subclause: 30C.6
Page: 90
Line: 13-19
CommentType (E, T, ER, or TR): E
Comment:
There appears to be an inconsistency between this MIB definition
and the attributes described in 30.7.4.2. The latter only has
current and expired as possible values here.
CommentEnd:

SuggestedRemedy:
Either remove values from the MIB definition here or add the new
values to 30.7.4.2 with some explanation to their meaning.
RemedyEnd:
Editor'sRecommendationStart:

The fault is with 30.7.4.2. This should reflect the possible states of the RX machine. Change the list to include:

current
expired
defaulted
initialize
lacpDisabled
portDisabled

and change the descriptive text to relate the enumerated values to the RX machine states.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 8
CommenterName: Graham Short
CommenterEmail: Graham.Short@tiuk.ti.com
CommenterPhone: +44-1604-663410
CommenterCompany: Texas Instruments
Clause: 43
Subclause: General
Page: many
Line: many
CommentType (E, T, ER, or TR): E
Comment:
There appear to be many instances that it looks as though the
"restart numbering" should have occurred for numbered bullets
(or in most cases "lettered" bullets) which have not occurred.

CommentEnd:

SuggestedRemedy:
Some of the ones I found are:

Page 111 line 24ff e)->h) should be a)->d)
Page 111 line 35ff i)->k) should be a)->c)
Page 112 line 13ff g)->i) should be a)->c)
Page 113 line 7ff e)->f) should be a)->b)
Page 113 line 29ff g)->h) should be a)->b)
Page 114 line 12ff c)->d) should be a)->b)
Page 115 line 15ff e)->f) should be a)->d)
Page 116 line 22ff c)->e) should be a)->c)
Page 118 line 10ff e)->f) should be a)->b)
Page 137 line 37ff d)->e) should be a)->b)
Page 139 line 25ff g)->j) should be a)->d)
Page 167 line 46ff c)->d) should be a)->b)
Page 169 line 46ff c)->e) should be a)->c)
RemedyEnd:
Editor'sRecommendationStart:

Reject. The draft is consistent with IEEE style guidelines as is. The style used ensures that references (e.g., 43.3 (c)) are unique and unambiguous.

Editor'sRecommendationEnd:
ResolutionStart:

ER. Also fix the one on p 114 (which is *incorrect*)

ResolutionEnd:

CommentNumber: 9
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: General
Page:
Line:
CommentType (E, T, ER, or TR): T
Comment:
This comment is by way of introduction to a rather extensive set of comments on clause 43, many of which are marked as TR (Technical Required). Although this is their strict designation, most of these comments are directed to aligning material in the draft that has lagged behind the extensive improvements made at the last meeting. The intent of these comments is thus not to introduce new technical material or ideas but bring the clauses commented on, many of which are introductory, into alignment with the progress made in the rest of the document. To make it

clear where I think a comment has a technical alignment purpose rather than a technical change throughout the document I have started the body of these comments with the label 'TA', sometimes with a (page, line reference).
CommentEnd:

SuggestedRemedy:
Adopt all comments marked TA so that the document is self consistent.
RemedyEnd:
Editor'sRecommendationStart:

AIP. Each comment should be considered on its merits, individually.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor’s Recommendation

ResolutionEnd:

CommentNumber: 10
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.1.2
Page: 100
Line: 6
CommentType (E, T, ER, or TR): E
Comment:

Groups of LANs is still unclear and could be more explicitly described. What I believe this is trying to describe is related to splitting an aggregator across multiple devices or systems. We aren't inventing a distributed MAC service interface where the MAC client can reside in separate nodes or bridges.

CommentEnd:

SuggestedRemedy:

Rename this to Multi-point Aggregations - the mechanisms in this clause only allow for a single aggregation to be established between two distinct systems.

RemedyEnd:
Editor'sRecommendationStart:

AIP. "Multipoint Aggregations - the mechanisms specified in this clause do not support aggregations among more than two systems."

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 11
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.2.1
Page: 102
Line: 51
CommentType (E, T, ER, or TR): E
Comment:

It is stated that the collector is free to select frames received from the aggregated ports in any order and pass them to its MAC Client in the order selected. This, perhaps, provides too much freedom.

CommentEnd:

SuggestedRemedy:

The collector is not required to reorder frames and should pass frames up to the MAC client specifically in the order received. If another, collector defined, order is selected or allowed, it should be done such that frame order is preserved.

RemedyEnd:
Editor'sRecommendationStart:

I believe that the existing text is correct, although it is admittedly "clear but not obvious". Consider changing "The collector is free to...", to:

"For any given port, the Collector is required to pass frames to the MAC Client in the order that they are received from that port. The Collector is otherwise free to select frames received from the aggregated ports in any order. Since there is no means..."

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 12
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.2.1
Page: 102
Line: 51
CommentType (E, T, ER, or TR): E
Comment:
"The collector is free to select frames received from the aggregated ports in any order and pass them to its MAC Client in the order selected"

It is not free to select frames in any order. It is free to select ports in any order. If the collector contained queues then it would be permitted to change the frame order from one link according to the current wording. (A hair-split I'll admit).
CommentEnd:

SuggestedRemedy:
Change to something like....

"The collector is free to service the aggregated ports in any order but will pass all frames received on a particular port to its MAC Client in the order they were received".

....but hopefully a little less verbose.
RemedyEnd:
Editor'sRecommendationStart:

See comment 11.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 13
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.2.3
Page: 103
Line: 52
CommentType (E, T, ER, or TR): E

Comment:
"Aggregator Parser" should be "Aggregator Parser/Multiplexer".
CommentEnd:

SuggestedRemedy:
"Aggregator Parser" should be "Aggregator Parser/Multiplexer".
RemedyEnd:
Editor'sRecommendationStart:

Reject. The collector sees only the Parser function of the Parser/Multiplexer, and this distinction was intentional.

Editor'sRecommendationEnd:
ResolutionStart:

Accept.

ResolutionEnd:

CommentNumber: 14
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.2.3.1.1
Page: 104
Line: 15
CommentType (E, T, ER, or TR): TR
Comment:
I think that defining CollectorMaxDelay as a default is insufficient, it really ought to be a maximum. An Actor may not know what CollectorMaxDelay has been set by administration at the Partner's end. If it has been changed to say 4 seconds then an Actor moving conversations after 1 second will cause frame misordering.

I think it would be better to exchange collector delay values via some new fields in the LACP. That way a Partner can explicitly tell an Actor how long it needs to wait before moving a conversation. If it declares say 1ms then this allows the Actor to move conversations hundreds of times per second instead of a maximum of 5 times at present (if it currently needs to use the Marker/Responder protocol).

If we did this then I wonder if the Marker/Response protocol can be removed?
CommentEnd:

SuggestedRemedy:
Change "default" to "maximum", or exchange collector delays via the LACP.

If the latter is adopted then consider if the Marker/Response protocol

is still required.
RemedyEnd:
Editor'sRecommendationStart:

Reject. We discussed this particular issue at the last meeting. At that time, we decided that a device could determine the value of its partner's CollectorMaxDelay through management, if it wanted to use a value other than the default.

In any case, knowing the precise value does not eliminate the need for the Marker protocol - for devices with a large value of CollectorMaxDelay, it will still be more efficient to flush rather than use a timeout.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Add the parameter to the LACPDU exchanges; make sure that the text adequately reflects the proper use of the parameter. Use a new TLV in the frame, properly aligned. Value is a constant, but implementation-dependent rather than one second.

ResolutionEnd:

CommentNumber: 15
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.2.3.1.4
Page: 104
Line: 46/47
CommentType (E, T, ER, or TR): E
Comment:

Missing closing brace on the Agg:MA_DATA.indication primitive

CommentEnd:

SuggestedRemedy:

Add a closing brace.

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 16
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.2.3.1.4
Page: 104
Line: Figure 43-3
CommentType (E, T, ER, or TR): T
Comment:

Strictly speaking, the Frame Collector diagram does not implement the functionality described in the text - namely, discarding any frames that cannot be delivered to the MAC Client within CollectorMaxDelay.

CommentEnd:

SuggestedRemedy:

Do we care about this discrepancy? If so, the diagram should be modified to reflect the specified discard action - this would involve adding a third state between the existing ones, with two exits, one to PASS TO PORT on the client accepting the frame, the other to WAIT FOR TRANSMIT on timer expiry.

RemedyEnd:
Editor'sRecommendationStart:

Reject. The subclause is correct as is.

First, I suspect the commenter meant “PASS TO MAC CLIENT” and “WAIT FOR RECEIVE” for the state transitions (this is the Collector, not the Distributor). However, the problem is that the current 802.3 architecture does not really provide for any such queueing or delay. This is why I “hand-waved” the issue in the text. From a conformance standpoint, there is no problem; the requirement is to conform to 43.2.3.1, which includes the text in question, in addition to Figure 43-3.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 17
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard

Clause: 43
Subclause: 43.2.4
Page: 105
Line: 33
CommentType (E, T, ER, or TR): E
Comment:

Frame flush mechanisms are not longer well described in place of the Marker protocol. Perhaps frame flush mechanisms includes both the Marker protocol and just waiting, but this is well described elsewhere.

CommentEnd:

SuggestedRemedy:

Either remove the note, or describe exactly what frame flush mechanism are. However, pointing the Annex 43A is a good idea.

RemedyEnd:
Editor'sRecommendationStart:

AIP. Re-word: "The subject of distribution algorithms and maintenance of frame ordering is discussed in Annex 43A."

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 18
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.2.4.1.3
Page: 106
Line: 19-23
CommentType (E, T, ER, or TR): TR
Comment:

The first of the Notes attached to Figure 43-4 expresses mandatory requirements.

CommentEnd:

SuggestedRemedy:

Elevate the first Note to the status of normative text. Remove the note number from the second Note.

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 19
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.2.7
Page: 107/108
Line:
CommentType (E, T, ER, or TR): TR
Comment:

Parsing of Marker frames should be conditional upon the DA being the Slow Protocols multicast address, in addition to the existing conditions. The specification of the protocol calls for the use of this address; it is unreasonable to expect implementations to accept marker frames transmitted to an arbitrary MAC address.

CommentEnd:

SuggestedRemedy:

In 43.2.7.1.1 add a new constant definition:

"Slow_Protocol_Address
The value of the Slow Protocols group MAC Address (see Table 43B-1)"

In Figure 43-5, change the exit condition from PARSE to PASS TO MARKER RESPONDER to:

"DA = Slow_Protocol_Address *
Length/Type = Slow_Protocol_Type *
subtype = Marker_Subtype *
TLV_type = Marker_Information"

In Figure 43-5, change the exit condition from PARSE to PASS TO MARKER

RECEIVER to:

"DA = Slow_Protocol_Address *
Length/Type = Slow_Protocol_Type *
subtype = Marker_Subtype *
TLV_type = Marker_Response_Information"

RemedyEnd:
Editor'sRecommendationStart:

Discuss. This exact text was *removed* from the last draft by specific instruction to the editor.

Editor'sRecommendationEnd:
ResolutionStart:

Accept. Add back the parsing of the MAC address to this, and the other SMs as needed.

ResolutionEnd:

CommentNumber: 20
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.2.7.1.4
Page: 108
Line: 5-36
CommentType (E, T, ER, or TR): TR
Comment:
The selection of the Marker Responder and the MAC Receiver from the PARSE state needs to be conditional on the MAC Address being the Slow Protocols Multicast. To do otherwise is to require that these functions effectively receive all frames promiscuously and filter on Ethertype alone at full line rate. This is a significant revision to 802 architecture and an unsupportable burden on most end stations and many switches.
CommentEnd:

SuggestedRemedy:
Add the destination address condition (Slow Protocols Multicast) to the transition to PASS TO MARKER RESPONDER and PASS TO MARKER RECEIVER.
RemedyEnd:
Editor'sRecommendationStart:

See comment 19.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 21
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.2.9.1.4
Page: 108
Line: 5-36
CommentType (E, T, ER, or TR): TR
Comment:
The selection of LACP from the PARSE state needs to be conditional on the MAC Address being the Slow Protocols Multicast. To do otherwise is to require that these functions effectively receive all frames promiscuously and filter on Ethertype alone at full line rate. This is a significant revision to 802 architecture and an unsupportable burden on most end stations and many switches.
CommentEnd:

SuggestedRemedy:
Add the destination address condition (Slow Protocols Multicast) to the transition to PASS TO LACP.
RemedyEnd:
Editor'sRecommendationStart:

See comment 19.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 22
CommenterName: Graham Short
CommenterEmail: Graham.Short@tiuk.ti.com
CommenterPhone: +44-1604-663410
CommenterCompany: Texas Instruments
Clause: 43
Subclause: 43.2.7.1.4 and 43.2.9.1.4
Page: 108 and 110
Line:
CommentType (E, T, ER, or TR): E
Comment:
Figures 43-5 and 43-6 used to have qualifiers for DA=Slow Protocol Multicast terms from the Parse boxes. These

appear to have gone in this version. I am not sure if this is because they don't need to be validated or if it is just an oversight
CommentEnd:

SuggestedRemedy:
Add DA terms into the diagrams if required.
RemedyEnd:
Editor'sRecommendationStart:

See comment 19.

Editor'sRecommendationEnd:
ResolutionStart:

Resolved as a technical comment.

ResolutionEnd:

CommentNumber: 23
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.2.7.1.4
Page: 108
Line: 18
CommentType (E, T, ER, or TR): TR
Comment:
The two "non-ELSE" exit conditions from the PARSE state were prefixed by the condition "DA = Slow_Protocol_Multicast *" in D1.0.

This was a necessary requirement for supporting existing devices that do not forward frames based upon the Length/Type field but instead forward frames based upon the destination address (only). These devices will forward frames carrying the Slow_Protocol_Multicast destination address to a CPU which can then verify the Length/Type field.

If a malformed frame which uses the Slow_Protocol_Type but doesn't use the Slow_Protocol_Multicast address is received on such a legacy device it will be forwarded to the MAC Client. The omission of "DA = Slow_Protocol_Multicast *" from these terms now changes this behaviour to make such legacy devices non-conformant with clause 43. They were always going to be non-conformant with 43B (Slow Protocols) but prior to this change could still claim conformance with clause 43 (Link Aggregation).

I tried to find this change in the Austin minutes and comment database, but couldn't locate it. Which comment was the change in response to?
CommentEnd:

SuggestedRemedy:

Prefix the non-ELSE exit conditions from the PARSE state with
"DA = Slow_Protocol_Multicast *" as in D1.0.
RemedyEnd:
Editor'sRecommendationStart:

See comment 19.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 24
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.2.7.1.4
Page: 108
Line: 23
CommentType (E, T, ER, or TR): E
Comment:

the FCS and status parameters in AggMuxN:MA_DATA.indication have been
dropped in some states but not others . They should be added back in
to be consistent.

CommentEnd:

SuggestedRemedy:

Make all AggMuxN:MA_DATA.indication calls the same.

RemedyEnd:
Editor'sRecommendationStart:

Reject. The FCS and status were intentionally left off the service calls to the Marker Responder and Marker Receiver, since those functions do not use those values for anything. They were passed to the Collector for the general client (e.g., a bridge) that may need them.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 25
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.2.7.1.4
Page: 108
Line: 26
CommentType (E, T, ER, or TR): T
Comment:
I wrote up this comment and then noted that Hadriel had submitted a similar comment on D1.0 which had been rejected (no reason given in comment database or minutes). I think that what should happen to unexpected Marker Responses if the Marker Receiver has not been implemented is still unclear. (Reject this comment if it is deemed out-of-order).

The Aggregator Parser state diagram (which is mandatory) shows it passing Marker Response frames to the Marker Receiver (which is optional). What happens to unexpectedly-received Marker Response frames if the Marker Receiver hasn't been implemented? Should they be quietly discarded, or should they be sent to the MAC Client because they are an unsupported Slow Protocol?
CommentEnd:

SuggestedRemedy:
Show Marker Response frames being quietly discarded (or passed to the MAC client) if the Marker Receiver hasn't been implemented. Add appropriate text to 43.2.6 (or thereabouts) too.
RemedyEnd:
Editor'sRecommendationStart:

AIP. There is a blanket requirement for Slow Protocols to pass unimplemented stuff up to the MAC Client. I.e., they are NOT quietly discarded. Clarify with a note and a forward reference.

Editor'sRecommendationEnd:
ResolutionStart:

Make the Marker Responder State and the exit condition “dotted line” (optional), Add a Note that this implies passing up Marker Responses to the Client if the Marker Receiver is not implemented.

ResolutionEnd:

CommentNumber: 26
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant

Clause: 43
Subclause: 43.2.9
Page: 109/110
Line:
CommentType (E, T, ER, or TR): TR
Comment:

Parsing of LACP frames should be conditional upon the DA being the Slow Protocols multicast address, in addition to the existing conditions. The specification of the protocol calls for the use of this address; it is unreasonable to expect implementations to accept LACP frames transmitted to an arbitrary MAC address.

CommentEnd:

SuggestedRemedy:

In 43.2.9.1.1 add a new constant definition:

"Slow_Protocol_Address
The value of the Slow Protocols group MAC Address (see Table 43B-1)"

In Figure 43-6, change the exit condition from PARSE to PASS TO LACP to:

"DA = Slow_Protocol_Address *
Length/Type = Slow_Protocol_Type *
subtype = LACP_Subtype"

RemedyEnd:
Editor'sRecommendationStart:

Discuss. This exact text was *removed* from the last draft by specific instruction to the editor. (Similar to comment 19 for Aggregator Parser.)

Editor'sRecommendationEnd:
ResolutionStart:

See comment 19

ResolutionEnd:

CommentNumber: 27
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.2.9.1.4
Page: 110
Line: 19
CommentType (E, T, ER, or TR): TR

Comment:
This is essentially the same comment as submitted for 43.2.7.1.4

The "non-ELSE" exit condition from the PARSE state was prefixed
by the condition "DA = Slow_Protocol_Multicast *" in D1.0.

This was a necessary requirement for supporting existing devices that do not
forward frames based upon the Length/Type field but instead forward
frames based upon the destination address (only). These devices will
forward frames carrying the Slow_Protocol_Multicast destination address
to a CPU which can then verify the Length/Type field.

If a malformed frame which uses the Slow_Protocol_Type but doesn't use the
Slow_Protocol_Multicast address is received on such a legacy device
it will be forwarded to the MAC Client. The omission of "DA = Slow_Protocol_Multicast *" from this term now
changes this behaviour to make such legacy devices non-conformant with
clause 43. They were always going to be non-conformant with 43B
(Slow Protocols) but prior to this change could still claim conformance
with clause 43 (Link Aggregation).

I tried to find this change in the Austin minutes and comment database,
but couldn't locate it. Which comment was the change in response to?
CommentEnd:

SuggestedRemedy:
Prefix the non-ELSE exit condition from the PARSE state with
"DA = Slow_Protocol_Multicast *" as in D1.0.
RemedyEnd:
Editor'sRecommendationStart:

See comment 26.

Editor'sRecommendationEnd:
ResolutionStart:

See comment 19

ResolutionEnd:

CommentNumber: 28
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.2.9.1.4
Page: 110
Line: 25
CommentType (E, T, ER, or TR): E
Comment:

The FCS and status parameters passed from MacN:MA_DATA.indication are not passed through CtrlMuxN:MA_DATA.indication and they are expected by the aggregator parser.

CommentEnd:

SuggestedRemedy:

Include all parameters in all xxx:MA_DATA.indication calls.

RemedyEnd:
Editor'sRecommendationStart:

AIP. Only the Aggregator needs the FCS and status.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 29
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3
Page: 111
Line:
CommentType (E, T, ER, or TR): TR
Comment:

The critical point that two ports that actually connect to each other on the same system should not be aggregated together even if they have the same key, but rather should be forced to be an individual link seems to have gone missing from this draft. It should appear in 43.3 as well as elsewhere
CommentEnd:

SuggestedRemedy:
Find a home for it and add it.
RemedyEnd:
Editor'sRecommendationStart:

AIP.

Add a note at the end of 43.3:

"System identification allows the detection of links that are connected in a loopback configuration (i.e., both ends of the same link connect to the same system); such links cannot be aggregated with any other link."

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Add note: "System identification allows the detection of links that are connected in a loopback configuration (i.e., both ends of the same link connect to the same system)."

Add text to the Selection Machine to indicate that a port cannot be aggregated with its partner. (Loopback case-don't allow a "reflector")

ResolutionEnd:

CommentNumber: 30
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.2
Page: 112
Line: 22
CommentType (E, T, ER, or TR): TR
Comment:

This subclause defines the System Identifier as a MAC address (a bit string or sequence of bits). In fact the System Identifier should be defined as a sequence of octets, in an analogous fashion to the Bridge Identifier defined in 802.1D-1998. This is key to getting the comparison operation correct so that the collating sequence of System Identifiers is as a normal human user would expect.

CommentEnd:

SuggestedRemedy:
Do it.
RemedyEnd:
Editor'sRecommendationStart:

Discuss. The intent of this section is to define the unique System MAC address. It could be expanded to also discuss the catenation with System Priority to provide the complete System Identifier.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Change to: "The globally unique identifier used to identify a system shall be a concatenation of a globally administered individual MAC

address (encoded as specified in xx), together with the System Priority. The MAC address chosen may be the individual MAC address associated with one of its ports" (think about better wording)

ResolutionEnd:

CommentNumber: 31
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.4
Page: 112
Line: 42-44
CommentType (E, T, ER, or TR): ER
Comment:

It is unnecessary and slightly confusing to propagate the duality of the terms "port identifier" and "port number" beyond the point where it is established that LACP uses a port number to identify a port. It has already been difficult enough distinguishing between the ifIndex as an identifier and the port number as an identifier.
CommentEnd:

SuggestedRemedy:
Replace
"43.3.4 Port identification

Each port is assigned a unique, globally administered individual MAC address. A port is also assigned an integer identifier (i.e., a port number) that is used by Link Aggregation Control to identify the port. Port identifiers shall be uniquely assigned within a system."

With
"43.3.4 Port identification

Each port is assigned a unique, globally administered individual MAC address. Link Aggregation Control also uses a port number to identify the port. Port numbers shall be uniquely assigned within a system."

RemedyEnd:
Editor'sRecommendationStart:

Accept. (I don't personally think there is much difference in the proposed wording; both it and the original text are correct. If the rewording makes the commenter happy, I don't have a problem with it.)

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 32
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.4
Page: 112
Line: 44
CommentType (E, T, ER, or TR): TR
Comment:

TA. A port number of zero should be explicitly disallowed so that zero can be used as a reserved value.
CommentEnd:

SuggestedRemedy:
At the end of 43.3.4 add the text.

"The port number of 0 should not be assigned."

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 33
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.5
Page: 112
Line: 48-
CommentType (E, T, ER, or TR): E
Comment:
This is the first time the reader encounters the term "key" in clause 43 apart from the terminology reference in 43.1.1 fourteen

pages prior. Reading through the document it seemed quite a shock to encounter this obviously important discussion which does not actually explain what keys are for.
CommentEnd:

SuggestedRemedy:
Modify line 42, page 112 through to the text of the paragraph starting at line 1 on page 113, so that the following existing text:

"43.3.5 Capability identification

A number of factors may affect the capabilities of a given port with respect to its ability to aggregate:

a) Its physical characteristics, such as data rate, duplexity, point-to-point or shared medium, etc.;
b) Configuration constraints established by the network administrator;
c) Factors related to higher layer use of the port (e.g. assignment of Network Layer addresses);
d) Characteristics or limitation of the port implementation itself.

In order to compare the capabilities of ports within a given system, and to allow the communication of such capabilities between systems, two Keys shall be associated with each port; an operational Key and an administrative Key."

is replaced with:

"43.3.5 Capability identification

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

a) The port's physical characteristics, such as data rate, duplexity, point-to-point or shared medium, etc.;
b) Configuration constraints established by the network administrator;
c) Use of the port by higher layer protocols (e.g. assignment of Network Layer addresses);
d) Characteristics or limitation of the port implementation itself.

Two Keys shall be associated with each port: an operational Key and an administrative Key."

RemedyEnd:

Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 34
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.3.5
Page: 113
Line: 7
CommentType (E, T, ER, or TR): E
Comment:
List starts with e)
CommentEnd:

SuggestedRemedy:
Start new list.
RemedyEnd:
Editor'sRecommendationStart:

Reject. See comment 8.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 35
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.3.5
Page: 113
Line: 7
CommentType (E, T, ER, or TR): E
Comment:
List starts with g)

CommentEnd:

SuggestedRemedy:
Start new list.
RemedyEnd:
Editor'sRecommendationStart:

Reject. See comment 8.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 36
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.3.5
Page: 113
Line: 32
CommentType (E, T, ER, or TR): E
Comment:
A link that is not a full duplex may not use the protocol and
therefore can not signal its capability. Another example is needed.
CommentEnd:

SuggestedRemedy:
Delete bullet h)
RemedyEnd:
Editor'sRecommendationStart:

Reject. This item does not imply the use of the LACP to signal its capabilities. The statement is correct as it stands; a system may determine that a link is not aggregatable through management or other out-of-band means.

Editor'sRecommendationEnd:
ResolutionStart:

Accept. Drop the second list element and subsume (g) into the text (for example...)

ResolutionEnd:

CommentNumber: 37
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084

CommenterCompany: 3Com
Clause: 43
Subclause: 43.3.5
Page: 113
Line: 39
CommentType (E, T, ER, or TR): TR
Comment:
I am concerned that we have not done sufficient failure analysis on the LAC Protocol.
For example with an AUI connection that does not give feedback as to the state of the receive signal in the event of only the receive cable being disconnected. The effect of this is that the two systems will not detect the link failure and may leave the port in the aggregation. The failure will result in one system loosing its receive data and be forced into the DEFAULTED state.
Now the problem occurs if the DEFAULTED information is the same as the previous operational information learnt from the neighbour. In this case the DEFAULTED system will expect the link to be still active and will continue to send traffic to its neighbour on the half failed link. The transmit path is still operational so it will maintain the far end in an operational state.
This failure can be fixed by ensuring that the DEFAULT data is changed if it is found that the learnt data from the neighbour matches the DEFAULT data.
There may be other cases like this one that will cause problems with real implementations.
Detailed remedy for this one is provided for 43.3.5, Page 113, Line: 39, 43.4.8, Page 130: Line 4, Fig 43-10 and 43.4.11, Page: 133, Line: 28.
CommentEnd:

SuggestedRemedy:
Add a paragraph to 43.3.5 at line 45 page 113:
To avoid problems with ports that are connected with AUI connections, and cannot detect failures in the receive path, the default and operational partner information must be different. Once the operational partner Key and operational partner System Identifier have been learnt from inbound Link Aggregation Control packets, the operational and default values for the partner are compared. Problems can occur if the default values of the partner Key and partner System Identifier, have the same values as the learnt operational partner Key and partner System Identifier. If this is found to be the case then the default values of the partner System Identifier of all ports using the same partner defaults shall be changed by a random amount to ensure that the default values are not the same as the operational ones. This change will ensure that both participants in the protocol detect and take corrective action in the event of a partial link failure.

Add a new function to the list in 43.4.8, Page 130, line 4:
modifyPartnerDefault
This function monitors the received Actor_System and Actor_Key and if the Actor_System and Actor_Key match the stored

Partner_Admin_Port_System_ID then the function modifies the values of the Partner_Admin_Port_System_ID for this and all ports that have the same value of Partner_Admin_Port_System_ID such that the received Actor_System is different to the admin default value.

Add a new function into Fig 43-10 section 43.4.11 Page 132.
Add to the CURRENT state after the update_NTT the new function: modifyPartnerDefault

Add a paragraph to 43.4.11, Page 133 at line 28:
When the state machine enters the CURRENT state the received Actor_System, Actor_Key are compared with the values of the Partner_Admin_Port_System_ID and the Partner_Admin_Port_Key. If the numerical values are the same then, to avoid problems with ports not running auto negotiation, the Partner_Admin_Port_System_ID for all ports with the same value must be changed by the same random amount such that received Actor_System is different to the operational Partner_Oper_Port_System_ID values.
RemedyEnd:
Editor'sRecommendationStart:

Discuss. It is not immediately clear that the proposed solution covers all the boundary conditions any better than the existing text. Perhaps a warning is needed rather than an automatic way of fixing the IDs. For example, the proposed procedure may not converge if two systems adopt the same, simplistic pseudo-random sequence.

Perhaps adopting a convention that default IDs are always taken from the locally administered address space would be useful.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Encode the value of the “Current/Expired/Default/Other” from the Receive Machine into the two reserved bits of the Actor/Partner state fields of the LACPDU.

ResolutionEnd:

CommentNumber: 38
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.3.6
Page: 114
Line: 36
CommentType (E, T, ER, or TR): E
Comment:
Confusing use of the word partners, this only occurs 5 times in the document and can easily be fixed.
Detailed remedies are provided for each of 43.3.6, Page 114, Line 36,

43.3.7 Page: 115 Line: 38, 43.3.7 Page: 115 Line: 41, 43.4.12 Page: 133 Line: 52, 43.6.1 Page: 146 Line: 3.
CommentEnd:

SuggestedRemedy:
43.3.6, Page 114, Line 36, Replace partners with participants in the Link Aggregation Control Protocol.
a) S and T are the respective System Identifiers of the participants in the Link Aggregation Control Protocol on the aggregated link(s);

43.3.7 Page: 115 Line: 38, Replace partners with participants.
The ability of LACP to signal to the participants that it considers a particular link to be Aggregatable or Individual results in optimization of the aggregation process in cases where the link is known by one or both parties to be Individual.

43.3.7 Page: 115 Line: 41, Replace partners with participants.
As there is no possibility of any extended negotiation between the two devices resulting in the participants agreeing to aggregate that link with any other link, both parties can take the necessary steps to form an Individual LAG for the link without further delay.

43.4.12 Page: 133 Line: 52, Replace partners with participants.
The Periodic Transmission machine establishes the desire of the participants to exchange periodic LACPDUs on the link in order to maintain an aggregation, and how often those periodic transmissions should occur.

43.6.1 Page: 146 Line: 3. Replace partners with participants in the Link Aggregation Control Protocol.
For a given pair of participants in the Link Aggregation Control Protocol, the system with the numerically lower value of System Aggregation Priority has the higher priority;

RemedyEnd:
Editor'sRecommendationStart:

AIP.
(1) Better to use "Actor and Partner"
(2) Better: "...signal to a potential aggregation partner..."
(3) Better to use "Actor and Partner"
(4) Better to use "Actor and Partner"
(5) Better to use "Actor and Partner", i.e. “For a given Actor and Partner...”

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 39
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.6
Page: 114
Line: 40
CommentType (E, T, ER, or TR): TR
Comment:

TA. This line states that "the system ID, operational Key and port identifiers are zero if not known". It is not completely clear if "not known" here means "can't be bothered with them" or "I've forgotten 'cause I'm lazy". On the assumption that it doesn't, it is no longer possible for these parameters to be unknown : the actor's own parameters are required to be always configured, the partner's parameters are as received in a LACPDU or are set to administrative defaults.

There should be a restriction preventing configuration of any of the actor's parameters to zero (somewhere), but it should be permissible to configure the partner default parameters to zero provided that the partner default parameters also identify the link as Individual.
CommentEnd:

SuggestedRemedy:
Delete line 40, page 114 (item d). State the above restrictions elsewhere.

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 40
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.6

Page: 114
Line: 3-44
CommentType (E, T, ER, or TR): ER
Comment:

TA. First explaining that a LAG ID is {SK, TL} and then explaining that it is not is a cumbersome explanation.
CommentEnd:

SuggestedRemedy:
Explain the LAG ID once. The following is suggested replacement text for the clause. It: adds the port number (not identifier) to the list of elements of the LAG ID; removes the first explanation of the symbols S, K, T, and L; removes the note but keeps the important comment that only one LAG can exist with the same SKTL and more than one member; fixes the rather haphazard use of labeled list items; and removes a few redundant words. The unchanged but relocated text is included below so that the effect of the proposed change is apparent.

"A Link Aggregation Group consists of either:

a) One or more Aggregatable links that terminate in the same pair of systems, and whose ports belong to the same Key Group in each system; or
b) An Individual link.

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

a) The System Identifier
b) The operational Key assigned to the ports in the LAG
c) The Port Number assigned to the lowest numbered Port in the LAG

A compound identifier formed from the System Identifier and Key values alone is sufficient to identify a LAG comprising Aggregatable links. This is not sufficient for a LAG comprised of a single Individual link where the partner System Identifier and operational Key may be unknown. Even if these are known there may be multiple Individual Links with the same System Identifier and operational Key combinations, and it is necessary to include port identifiers to provide unique LAG IDs.

If:

a) S and T are System Identifiers
b) K and L are the operational keys assigned to a LAG by S and T respectively
c) P and Q are the Port Numbers of the lowest numbered ports in the LAG at S and T

the general form of the unique LAG ID is {SKP, TLQ}.

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

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 41
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.6
Page: 114
Line: 3-44
CommentType (E, T, ER, or TR): TR
Comment:

At present it is not possible to compare LAG IDs to see whether two system with the same LAG ID should be in the same LAG. To do so requires knowledge of whether the LAG contains an Individual link or not.
CommentEnd:

Suggested Remedy:

Establish the convention that P and Q are ALWAYS BOTH zero if the LAG is Aggregatable, and NEVER BOTH zero if it is not. This seems an easier fix than adding an I or A prefix to the LAG ID. It does require that zero is never configured as a port number for actor information.

Note : Since both actor and partner information are always available (although in some cases the partner information will be the default), the rest of the LAG ID should always be known.

If the prior comment on 43.3.6 simplifying the description of the LAG ID is accepted, change the text describing the construction of the LAG ID (item c) to:

"c) The Port Number if the link is identified as an Individual link"

and the text describing P and Q to:

"c) P and Q are the Port Numbers of the ports attached if the LAG comprises a single Individual Link and zero if the LAG comprises one or more Aggregatable links."
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 42
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.6
Page: 114
Line:
CommentType (E, T, ER, or TR): TR
Comment:
This clauses refers to the numerical comparison of System Identifiers. But System Identifiers, as defined in 43.3.2, are not numbers. They are MAC Addresses, and MAC Addresses are bit strings. The procedure for the comparison of System Identifiers needs to be made explicit. To facilitate that comparison System Identifiers should be defined as octet sequences.
CommentEnd:

SuggestedRemedy:
Define the comparison process, here or in 43.3.2, using the same procedure as described in 802.1D-1998 Clause 9.2.5.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Pending earlier resolution.

ResolutionEnd:

CommentNumber: 43
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.6
Page: 114
Line:
CommentType (E, T, ER, or TR): TR
Comment:
It is unclear as to whether the system and port priority components form part of the LAG ID or not.
CommentEnd:

SuggestedRemedy:
State whether or not they are.
RemedyEnd:
Editor'sRecommendationStart:

Accept. (They are part of LAG ID.) This should fall out of incorporating the priority elements into 43.3.2 and 43.3.4, per earlier comments.

Editor'sRecommendationEnd:
ResolutionStart:

ER Done as part of earlier comment

ResolutionEnd:

CommentNumber: 44
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.6
Page: 114
Line:
CommentType (E, T, ER, or TR): TR
Comment:
The LAG ID should have a clear human readable representation. Examples should be provided, so that it is clear to the reader and to us what we have done.
CommentEnd:

SuggestedRemedy:

Define a representation including the System Identifier derived from the MAC address in 'canonical form' and with sufficient punctuation to make the components of the LAG ID clear.

The representations should include examples of LAG Ids.

RemedyEnd:
Editor'sRecommendationStart:

AIP. Prepare text and example based on the slide presented at the meeting. (Use comma notation and give an example of the FULL (both partners) LAGID syntax.

Editor'sRecommendationEnd:
ResolutionStart:

ER. Proposed syntax is to use “*” for the punctuation delimiter for the priority field prepended to the MAC address. Mick to prepare an example offline.

ResolutionEnd:

CommentNumber: 45
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.7
Page:
Line:
CommentType (E, T, ER, or TR): TR
Comment:
TA. This subclause confuses the selection of a Link Aggregation Group with the agreement that needs to be reached between the communicating systems before the link is used. In LACP these correspond to the two separate ideas of Selection and Matching. This comment addresses only the current contents 43.3.7 in detail. In addition subclause 43.3.7 begins by describing the distinction between Individual and Aggregatable links that has just be presented in 43.3.6. Since we have agreed that an aggregate of one link is a valid aggregate and needs to be attached to an aggregator this repetition is unnecessary. The only points that need to be made in 43.3.7 concern the precise conditions for LAG selection and the utility of signaling Individual to accelerate that selection. The LAG ID has already been described in 43.3.6. The subclause also needs updating to reflect the change made in the last draft to include administrative default values for the partner system. Thus it is not the case, as may be inferred from the current text, that the lack of information from the partner system implies that the link is not Aggregatable.

CommentEnd:
SuggestedRemedy:
Replace the entire subclause 43.3.7 with the following text:

"43.3.7 Selecting a Link Aggregation Group

Each port, and its attached link, is selected for membership in the Link Aggregation Group uniquely identified by the LAG ID comprised 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, otherwise locally configured administrative defaults are assumed for the remote 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."

Add a new subclause 43.3.8 (and renumber following subclauses) with the following text:

"43.3.8 Agreeing 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 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."

RemedyEnd:
Editor'sRecommendationStart:

AIP. The editor should wordsmith the proposed text.

Editor'sRecommendationEnd:
ResolutionStart:

ER Make sure that nothing is LOST from this section by the change (other changes need to be kept).

ResolutionEnd:

CommentNumber: 46
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.3.7
Page: 115
Line: 15
CommentType (E, T, ER, or TR): E
Comment:
List starts with c)
CommentEnd:

SuggestedRemedy:
Start new list.
RemedyEnd:
Editor'sRecommendationStart:

Reject. See comment 8.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 47
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.3.8
Page: 116
Line: 23
CommentType (E, T, ER, or TR): E
Comment:
Item d) has a PICS entry (ALA2) suggesting there should be
a "shall" somewhere in this sentence.
CommentEnd:

SuggestedRemedy:
Add a "shall" (or remove the PICS).
RemedyEnd:
Editor'sRecommendationStart:

Accept. Add the “shall”.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 48
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.3.8
Page: 116
Line: 21
CommentType (E, T, ER, or TR): E
Comment:
List starts with c)
CommentEnd:

SuggestedRemedy:
Start new list.
RemedyEnd:
Editor'sRecommendationStart:

Reject. See comment 8.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 49
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.3.8
Page: 116
Line: 25
CommentType (E, T, ER, or TR): TR
Comment:
The "shall" in this list item requires Collection to occur before Distribution begins. The note at the top of page 117 says that this is not a requirement. These appear contradictory.

CommentEnd:

SuggestedRemedy:
Roll the note into the list item and make clear exactly what
is required.
RemedyEnd:
Editor'sRecommendationStart:

AIP. Collection should occur at or before the time that Distribution begins, to allow for implementations that cannot separate the behaviors.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. line 20 change “must be met” to “apply”
Change wording of c) d) and e) to be informative (remove shall requirements)
Change line 25-6 to “(i.e., this allows Collection to occur with respect to a given link before Distribution begins).

ResolutionEnd:

CommentNumber: 50
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.3.8
Page:
Line:
CommentType (E, T, ER, or TR): TR
Comment:
This subclause confuses the idea of the attachment of a link to an Aggregator with the process of enabling and disabling Distribution and Collection. The notion of attachment to an Aggregator is an abstraction for the organization of local resources in preparation for transmission and reception.

Once a link is 'attached' to the correct Aggregator, any attempt to transmit frames on the link should be for frames intended for transmission by the logical MAC interface attached to the agreed upon LAG. Frames received by the peer port might be delivered to the MAC client as if they had been transmitted by a peer MAC client attached to the agreed upon LAG. However attachment to an Aggregator can happen when the Link Aggregation Group is selected (43.3.7) and before the LAG ID is agreed (suggested new subclause after 43.3.7).It is therefore inappropriate to begin this subclause (43.3.8) with a discussion of the frame ordering issue.

The following technical issues need to be discussed, in this

order:
Attaching a link to an Aggregator
Signaling readiness to transfer user data
Enabling Collection and Distribution

The first of these, Attaching a link to an Aggregator, needs to cover the following topics:
Finding a compatible Aggregator. (as currently described in 43.3.8)
Avoiding hysteresis (as currently described in 43.3.8)
Links that are not successful candidates for aggregation (as currently described in 43.3.8)

Additional problems in this area that the Remedy suggested below also addresses are:
If the port's operational Key is to be changed from its administrative value then the Aggregator's Key must also be changeable, otherwise there might not be a compatible Aggregator. There appear to be two ways to fix this. One is to allow such a change of Aggregator key (and this appears to be supported by the rest of the document). The other is to select an Aggregator based on the comparison of administrative keys, and is the subject of a later comment. For the present the Remedy suggests the former, so as not to bundle too many issues into this single comment.
An unnecessary distinction is being made between aggregatable links and individual links. This distinction can be entirely contained within subclause (43.3.7). See proposed remedy text in an earlier comment. Compatibility of links attached to an Aggregator can be entirely assured by them having chosen the same Link Aggregation Group.
The description of the reasons for applying hysteresis should include the concept that, if additional links select the same LAG within a short period of time, the Aggregator that is appropriate may change, particularly if some form of deterministic selection is being applied.

The second, Signaling readiness to transfer user data, needs to cover the following:
Notifying the peer that you have attached to an aggregator, and that the aggregator is compatible with the agreed upon LAG ID (not currently described in 43.3.8)
Ensuring that there are no outstanding frames on the link (beginning paragraph of current 43.3.8). Note that this condition should be met before notifying the peer that everything is ticketty boo.

Additional problems that the Remedy suggested below also addresses are:
c. The quality of service aspects to be considered when moving links between Aggregators include misdelivery as well as order preservation and preventing duplication. For this reason the

concern has to be whether there are frames on a link transmitted while the link was a member of a prior Link Aggregation Group, not just while the link was attached to a prior Aggregator. LACP deals with the change of LAG case correctly.

The third, Enabling Collection and Distribution, needs to cover the following:
Receiving notification from the peer system that it has attached to an Aggregator compatible with the agreed upon LAG ID, and that there are no frames in transit from that peer to the actor (not currently described in 43.3.8)
Turning the Collector on (as currently described).
Turning the Distributor on (as currently described).
Reflecting Distributor and Collector enabled / disabled states into Aggregator states (as currently described).
Explaining what happens when the implementation cannot separately enable collection and distribution.

CommentEnd:

SuggestedRemedy:
Replace the entire subclause 43.3.8 with the following text. Note that I have numbered new subclauses starting with 43.3.9, on the assumption that 43.3.7 would be split as suggested in a previous comment. I have used as much of the existing text as possible in suggesting this Remedy but the cutting and pasting is a little elaborate to describe precisely. In summary the suggested text comprises (in order) text from the current first paragraph with the removal of the discussion of frames in transit, text from the second paragraph simplified to not treat individual links as a special case, the addition of a description of the optional procedures used when there are many compatible Aggregators, text moved from the second paragraph to discuss the case when no compatible Aggregator exists, the existing NOTE unchanged, the text of the existing fifth paragraph unchanged, the text of the existing fourth paragraph with minor grammatical simplifications and corrected to discuss attachment of ports to an Aggregator specifically rather than the aggregation of links in general.

"43.3.9 Attaching a link to an Aggregator

Once a link has selected a Link Aggregation Group, Link Aggregation Control can attach the link to a compatible Aggregator. An Aggregator is compatible if:
a) The Aggregator's operational Key matches the 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 to ensure deterministic

behavior, i.e. independence from the order in which Aggregators become available, or to maximize availability of the aggregation to a MAC 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 service 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-aggregatable) are enabled to operate as individual 802.3 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."

Add a new subclause 43.3.10 with the following text (the heading could do with some improvement). The signal referred to is of course the In_Sync signal of LACP. In summary the suggested text is largely that displaced from the first paragraph of the existing 43.3.8, with clarification of the conditions under which each peer LAC entity is ready to transfer data, and explicit mention of that signaling.

"43.3.10 Signaling readiness to transfer user data

Once a link has been attached to an Aggregator (43.3.9) compatible with the agreed Link Aggregation Group (43.3.8) , each Link Aggregation Control entity signals to its peer readiness to transfer user data to and from the Aggregator's MAC 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 prevention of frame duplication. Link Aggregation Control will not signal readiness until it is certain that there are no frames in transit on the link that were transmitted while the link was a member of a previous Link Aggregation Group. This may involve the use of an explicit Marker protocol that ensures that no frames remain to be received at either end of the link before reconfiguration takes place. The operation of the Marker protocol is described in 43.5. The decision as to when, or if, the Marker Protocol is used is entirely dependent upon the nature of the distribution algorithm that is employed."

Add a new subclause 43.3.11 with the following text. In summary the suggested text takes the sixth paragraph of the existing 43.3.8, clarifies when the process of enabling the link can proceed in terms of the proposed 43.3.10, moves the asides as to what happens to the Aggregators receive and transmit state to a final paragraph, and cuts out some needless elegant variation where the Receive state refers to 'at least one port' and the Transmit state to 'at least one link'.

"43.3.11 Enabling Collection and Distribution

Every Aggregator can enable or disable Collection and Distribution of frames for each port that is attached to the Aggregator. Initially both Collection and Distribution are disabled. Once the Link Aggregation Control entity is ready to transfer user data using the link, and its peer entity has also signaled readiness, the process of enabling the link can proceed. The Collector is enabled, thus preparing it to receive frames sent over the link by the remote Aggregator's Distributor, and that fact is communicated to the remote Link Aggregation Control entity. Once information received indicates that the remote Aggregator's Collector is enabled, the Distributor is also enabled.

NOTE - This description assumes that the implementation is capable of controlling the state of the transmit and receive functions of the MAC independently. In implementation where this is not possible, the transmit and receive functions are enabled or disabled together. The manner in which this is achieved is detailed in the description of the Mux state machine (see 43.4.14).

If at least one port in the Link Aggregation Group has its Collector enabled, then the Receive state of the corresponding Aggregator will also be Enabled. If at least one port in the Link Aggregation Group has its Distributor enabled, then the Transmit state of the corresponding Aggregator will also be Enabled."

RemedyEnd:
Editor'sRecommendationStart:

AIP. The editor should wordsmith the proposed text.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 51
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.3.12
Page: 118
Line: 21
CommentType (E, T, ER, or TR): E
Comment:
List starts with e)
CommentEnd:

SuggestedRemedy:
Start new list.
RemedyEnd:
Editor'sRecommendationStart:

Reject. See comment 8.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 52
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.1 LACP design principle
Page: 118
Line: 30-
CommentType (E, T, ER, or TR): TR

Comment:
These are not just design principles but a rag bag of observations.

CommentEnd:
SuggestedRemedy:
Put design principles in this subclause.
RemedyEnd:
Editor'sRecommendationStart:

Discuss specifics.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Change section title to “LACP Design Elements”, and the first line to “the following considerations were taken into account in the development...” Wordsmith as appropriate.

ResolutionEnd:

CommentNumber: 52a
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.4.1
Page: 118
Line: 37
CommentType (E, T, ER, or TR): ER
Comment:

The description of Active or Passive participation is not entirely clear or consistent. The use of the terms "willingness to participate in the protocol" implies a Passive participant will never send LACPDUs and certainly not on its own initiative. However, the MUX machine will likely cause NTT to assert at initialization time, so at least one LACPDU will likely be sent regardless of Passive or Active. The description of Periodic Transmission Machine on page 122, line 37 describes how Active and Passive participation impact the transmission of periodic LACPDUs for the purpose of keeping the aggregation alive. This description seems more accurate.

CommentEnd:

SuggestedRemedy:

Very clearly explain the operation of Active and Passive with respect to the transmission of LACPDUs. If the true meaning is, "don't speak until spoken too" then modify the MUX machine such that this will be

true. If the intention is to allow an aggregation to go idle and force default parameters to be loaded, then describe this. Otherwise, change the description of Active and Passive to describe that it only impacts the process of initiating and sending periodic transmissions, and that a Passive participant may send and even initiate LACPDUs.

RemedyEnd:
Editor'sRecommendationStart:

Discuss.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. If both Actor and Partner are passive, clear NTT (preventing all transmissions) in the Transmit SM. Clarify 43.4.1 (c) to show that Passive means a preference for not transmitting unless necessary, rather than a “willingness to participate”.
On p122 L40, delete the phrase after the comma. Add a statement to line 47 (e) that when Actor and Partner are both passive, then no transmissions occur.

ResolutionEnd:

CommentNumber: 53
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.2
Page:
Line:
CommentType (E, T, ER, or TR): TR
Comment:

The encoding of items described as being "binary numbers" in this subclause will lead to a great deal of confusion and must be changed. If for some perverse reason it has to stay as it is a clear warning in no less than 36 point type and containing no fewer than 100 alphabetic characters must be added to the document. The encoding of "numbers" should be spelled out precisely, as should the encoding of MAC addresses, and should be the same as that presented in 802.1D-1998. In addition the ordering of some of the fields should be changed to support the required comparisons without further manipulation.

CommentEnd:

SuggestedRemedy:
Replace item c) of 43.4.2.1 with:
"c) When consecutive octets are used to represent a binary number, the octet transmitted first has the more significant

value."

Add an item d) as follows:

"d) When consecutive octets are used to represent a System Identifier, the least significant bit of the first octet is assigned the value of the first bit of the MAC address used to derive the System Identifier, the next most significant bit the value of the second bit of the MAC address, and so on through the eighth bit. Similarly the least significant through most significant bits of the second octet are assigned the value of the ninth through seventeenth bits of the MAC address, and so on for all the octets of the System Identifier."

In 43.4.2.2 and Figure 43-7 reverse the order in which the following pairs of parameters appear: Actor_System and Actor_System_Priority; Actor_Port and Actor_Port_Priority; Partner_System and Partner_Port_Priority.

RemedyEnd:
Editor'sRecommendationStart:

Discuss. Personally, I believe that designating a value as a “binary number” is much more useful than calling it an octet. Numbers have numerical value and take arithmetic operations and comparisons better than generic “octets”, but I am open to TF input. Since the usage of System and Port IDs is comparable to 802.1D, I don’t have a big problem with a consistent definition of the syntax.

The specific wording is clumsy, however. It suffices to say that, for multi-octet fields, the first octet is the most significant, the second is the next most significant, etc. The bit significance is already defined for an octet, as is the transmission bit order of 802.3.

Accept the change of parameter order, to align the 6-octet fields on word boundaries.

Editor'sRecommendationEnd:
ResolutionStart:

Accept.

ResolutionEnd:

CommentNumber: 54
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.2.2
Page:
Line:
CommentType (E, T, ER, or TR): TR

Comment:

The Actor_Port_Priority and Partner_Port_Priority fields should be of the same size as the Actor_Port and Partner_Port fields to allow for easy arbitrary ordering of all the ports on a system.

A PAD octet should be added following the Actor_State and Partner_State fields to maintain even octet alignment.
CommentEnd:

SuggestedRemedy:
Change the port priority field sizes to two octets and add a PAD octet after the state fields.
RemedyEnd:
Editor'sRecommendationStart:

Discuss.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Pad to 4-byte alignment of TLVs.

ResolutionEnd:

CommentNumber: 55
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.4.2.2
Page: 119
Line: 24
CommentType (E, T, ER, or TR): ER
Comment:

The use of Slow_Protocols_Multicast address in bullet a) does not align with terminology in 43B.

CommentEnd:

SuggestedRemedy:

Replace "Slow_Protocols_Multicast address" with "Slow Protocols group MAC Address".

RemedyEnd:
Editor'sRecommendationStart:

AIP. Consider using “Slow_Protocols_Multicast” consistently, rather than “group MAC address”.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 56
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.4.2.2
Page: 119
Line: 28
CommentType (E, T, ER, or TR): ER
Comment:

The use of Slow_Protocol type field in bullet c) does not align with terminology in 43B.

CommentEnd:

SuggestedRemedy:

Replace "Slow_Protocol type field" with "Slow Protocols Ethernet Type"

RemedyEnd:
Editor'sRecommendationStart:

AIP. Correct 43B to reflect the term “Slow_Protocol type field”. As this is an Ethernet-only standard, this is the correct terminology (per Clause 3).

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 57
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.4.3
Page: 123
Line: 5
CommentType (E, T, ER, or TR): E

Comment:

The arrows seem incorrect in this diagram. The Actor State information either needs to go into all states or only the TX state.

CommentEnd:

SuggestedRemedy:

Show an arrow into all states that read Actor state information as well, or show that the flow of Actor state information is towards TX.

RemedyEnd:
Editor'sRecommendationStart:

Reject. I believe it is OK as is. The only machines that need to know about the state of *other* machines (other than itself) are TX and RX; TX because it needs to transmit stuff out, and RX because it needs to compare actual Actor's state versus Partner's view of Actor's state.

Editor'sRecommendationEnd:
ResolutionStart:

Withdrawn.

ResolutionEnd:

CommentNumber: 58
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.3
Page: 123
Line: 18
CommentType (E, T, ER, or TR): T
Comment:
The diagram does not show the inter-relationship between this port and the other ports in the aggregation. The ready variable is used in the selection logic to indicate that the wait_while timers of all other ports in the same aggregation have expired. This input to the selection machine from other ports in the same aggregation should be indicated on this top level diagram because it is a very significant input to the selection machine.
CommentEnd:
SuggestedRemedy:
43.4.3 Page 123 line 18 Fig 43-9
Add several inputs ready_M into the selection machine to represent the inter-relationship between this port and other ports in the same

aggregation.

43.4.7 Page 129 line 17
Change the description of ready: replace the words
and the wait_while timers for all ports that are ready to attach to the Aggregator .. have expired
with
and the value of ready_N to be TRUE for all the ports that are attached to the Aggregator

43.4.7 Page 129 line 16
Add a definition for ready_N:
ready_N
A value of TRUE indicates that the Selection logic is ready to be attached to the aggregator and its wait_while timer has expired. A value of FALSE indicates that either the port is not ready to be attached to the aggregator or its wait_while timer is still running. This variable is used by other ports to determine the value of the ports ready variable.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 59
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.4
Page: 124
Line: 21
CommentType (E, T, ER, or TR): E
Comment:
Replace Churn_Detection__Time with Churn_Detection_Time
CommentEnd:

SuggestedRemedy:
Replace Churn_Detection__Time with Churn_Detection_Time
RemedyEnd:
Editor'sRecommendationStart:

Accept, and get a referral to Mr. Balmer’s optometrist.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 60
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.5
Page: 125
Line: 29
CommentType (E, T, ER, or TR): T
Comment:
Unknown value is not defined.
CommentEnd:

SuggestedRemedy:
Add 'Zero indicates "unknown"'
RemedyEnd:
Editor'sRecommendationStart:

Align with Mick Seaman's earlier comments on unknown values.

Editor'sRecommendationEnd:
ResolutionStart:

Accept.

ResolutionEnd:

CommentNumber: 61
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.6
Page: 126
Line: 11
CommentType (E, T, ER, or TR): T
Comment:
The value zero for unknown System_ID values needs to be defined consistently across the spec. In particular since it is defined in the definition of Partner_System_ID for the aggregator in section 43.4.5, it also is needed in the port information in section 43.4.6.

Also in Subclause: 43.3.6 Page: 114 Line: 40
The bullet d) implies that the SystemID may take on either the identifiers of the participants or a default value or zero if no default is specified. The Key may be either a real value or a default value (see for example aAggPartnerOperKey or Partner_Oper_Key definitions). The bullet suggests that two of these three options are true for both. It needs changing.
CommentEnd:

SuggestedRemedy:
Add to definition of Partner_Admin_System_ID 43.4.6 page 127 line 13
Add the words
0x00-00-00-00-00-00 indicates unknown
Add to definition of Partner_Admin_Key 43.4.6 page 127 line 31
Add the words
Zero indicates unknown
Change the bullet d) on page 114 as follows:
d) The system ID, operational Key and port identifiers for the partner will default to an administratively set value or zero if not known.
RemedyEnd:
Editor'sRecommendationStart:

Discuss. See comment 60.

Editor'sRecommendationEnd:
ResolutionStart:

Reject. The Partner values are never “unknown”; if not imparted by the protocol, they go to the Administrative default.

Clarify definition of Partner Admin parameters to indicate that they are the *Actor’s* Administative value for the Partner’s parameter.

ResolutionEnd:

CommentNumber: 62
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.4.6
Page: 126
Line: 31
CommentType (E, T, ER, or TR): E
Comment:

There is only one remote system - known as the Partner.

CommentEnd:

SuggestedRemedy:

Change "remote system(s)" to "Partner".

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 63
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.6
Page: 128
Line: 30
CommentType (E, T, ER, or TR): TR
Comment:
port_enabled is used also to indicate link state. This must be clearly
specified in the spec to avoid confusion.
CommentEnd:

SuggestedRemedy:
Change the definition of port_enabled as follows:
A signal indicating that the physical layer has indicated that the link
has been established and the port is operable.
RemedyEnd:
Editor'sRecommendationStart:

AIP. Whether a given port makes such an explicit signal available, and whether it is present on a Physical Layer interface is implementation-dependent. Discuss rewording.

Editor'sRecommendationEnd:
ResolutionStart:

Accept.

ResolutionEnd:

CommentNumber: 64
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com

CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.7
Page: 128
Line: 47
CommentType (E, T, ER, or TR): E
Comment:
This "shall" doesn't appear to be in the PICS.
CommentEnd:

SuggestedRemedy:
Add to PICS or remove "shall" from this sentence.
RemedyEnd:
Editor'sRecommendationStart:

Add to PICS.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 65
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.4.7
Page: 129
Line: 32
CommentType (E, T, ER, or TR): T
Comment:

The description of how Matched is set TRUE does not match what is described in the choose_Matched function or how it is assigned in the statemachine. The description says matched is set TRUE when the receive state machine is using default values for the partner's operational state. However, it is possible to be in both the PORT_DISABLED and EXPIRED states with the default values loaded, but matched set to FALSE. Also the choose_Matched function says Match is set TRUE if the value of Actor_State.Aggregation in the received PDU is set to FALSE, but this is not indicated in the description.

CommentEnd:

SuggestedRemedy:

Matched should only be set to TRUE if the Actor and Partner parameters correspond, or if the link is individual. The default parameters should be responsible for assuring the Actor and Partner parameters correspond. choose_Matched should be called in the PORT_DISABLED and DEFAULTED states rather than being explicitly assigned.

RemedyEnd:
Editor'sRecommendationStart:

Reject. In the EXPIRED state, you are not matched; similarly if the Port is disabled. I don't believe that this is broken.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Add an additional qualifier at p129 l32, “or the partner is Individual”

ResolutionEnd:

CommentNumber: 66
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.8
Page: 130
Line: 3
CommentType (E, T, ER, or TR): E
Comment:
Wrong name for Partner_Admin_Port_State and Partner_Oper_Port_State _Port_ missing on line 3, 9, 11, 17, 20, 29!!
CommentEnd:

SuggestedRemedy:
On lines 3, 9, 11, 17, 20, 29 use Partner_Admin_Port_State and Partner_Oper_Port_State in these descriptions.

Note that line 1, 26 are ok because this is the LACPDU name!
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 67
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.8
Page: 130
Line: 9,11,17,20,29,37,40,52
CommentType (E, T, ER, or TR): E
Comment:
Change "Partner_Admin_State" to "Partner_Admin_Port_State"
Change "Partner_Oper_State" to "Partner_Oper_Port_State"
CommentEnd:

SuggestedRemedy:
Change "Partner_Admin_State" to "Partner_Admin_Port_State"
Change "Partner_Oper_State" to "Partner_Oper_Port_State"
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 68
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.8
Page: 130
Line: 34
CommentType (E, T, ER, or TR): TR
Comment:
update_Default_Selected can only be called when no LACP frame has been received. At which point the default port parameters are transferred to the operational port parameters using function recordDefaultPDU. The default port parameters will now match the operational port parameters, because they have been copied. The description of update_Default_Selected function has to be changed to reflect that this fact, it does not depend upon a received PDU as presently stated. The comparision between the new values of the operational port parameters is

performed by the selection machine which may decide that the port is to remain selected even though the defaults are now being used rather than any received information. The receive machine can not make this decision.
Same argument for the update_LACPDisabled_Selected function. Since recordLACPDisabledPDU copies the information this function also is not required since it will not alter the value of the Selected variable.
Only the selection machine has the knowledge about the validity of the new default operational state information.
CommentEnd:

SuggestedRemedy:
Delete the update_Default_Selected function from 43.4.8 page 130 line 34 and from Fig 43-10
Delete the update_LACPDisabled_Selected function from 43.4.8 page 130 line 44 and from Fig 43-10
RemedyEnd:
Editor'sRecommendationStart:

Reject. Change the description to reflect the fact that the comparison is with default values, not those received in a LACPDU; however, the rationale for both of these operations is no different from the rationale for update_Selected; i.e., to force de-selection if new information is plugged into the state machines & it conflicts with the current selection.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 69
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.4.8
Page: 130
Line: 42
CommentType (E, T, ER, or TR): TR
Comment:

The update_Default_Selected function indicates that Selected remains unchanged if the Partner Operational and Partner Administrative values are the same, however, this will prohibit successful aggregation for a locally administered situation starting from the INITIALIZATION state when the partner is not running LACP. Selected is set FALSE in INITIALIZATION when the default parameters are loaded. It remains unchanged when entering EXPIRED. If no LACP frames are ever received, the machine transitions to DEFAULTED where the update_Default_Selected

does not change the current value of FALSE to TRUE. Without Selected set TRUE the selection statemachine will never leave the DETACHED state.

CommentEnd:

SuggestedRemedy:

update_Default_Selected should always set Selected TRUE if the Partner Operational and Partner Administrative values are the same.

RemedyEnd:
Editor'sRecommendationStart:

Reject. The Selection machine will correctly do its thing simply by running the selection logic. That is how selection leaves DETACHED. The RX machine never sets selected=TRUE, and should not do so. That would be equivalent to importing the selection machine into the RX machine, which is not desirable.

Editor'sRecommendationEnd:
ResolutionStart:

AIP.

(1) Some additional words are be needed to clarify the relationship between machines. In particular, it should clearly state that the RX machine will only set Selected=FALSE and the Selection machine will set Selecte=TRUE or STANDBY. Perhaps the 'update_selected' function should be called 'verify_selected' or 'determine_if_still_selected_or_standby' or something like that.

It would also be helpful to describe how the Matched variable is managed and what is the relationship between RX, Selection and Mux with respect to setting this variable and sending signals.

(2)The description of the Matched variable should be changed to indicate the other situations where it is set FALSE by the RX machine. There are cases where the appropriate parameters are loaded to have Matched=TRUE, but certain states are driving it FALSE.

(3) Add to definition of “selected” a reference to all of the SMs that affect its value. Clarify that Selection logic is continuously executing.

ResolutionEnd:

CommentNumber: 70
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43

Subclause: 43.4.7
Page: 130
Line: 50
CommentType (E, T, ER, or TR): E
Comment:
Its only a minor point but all these comparisons are all "equality" comparisons with the exception of comparing "Individual" with "Partner_Oper_Port_State.Aggregation". In this case the values compared are of opposite sense.
CommentEnd:

SuggestedRemedy:
Perhaps make this clear, or change one variable's sense to match the other, or do nothing. (Editorial discretion).
RemedyEnd:
Editor'sRecommendationStart:

The editor will be discreet.

Editor'sRecommendationEnd:
ResolutionStart:

Resolved by technical comment resolution.

ResolutionEnd:

CommentNumber: 71
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.9
Page: 131
Line: 28
CommentType (E, T, ER, or TR): TR
Comment:
The protocol as specified is unresilient to failures. With only one lost PDU the aggregation could be forced to toggle due to Matched=FALSE in the Expired state, resulting in out of sync in the Mux. In addition both the participants (actor and partner) will have to synchronise so that the transmitted LACPDU arrives in time to prevent the receive machine timeout. This is a direct result of starting the current_while_timer each time the receive engine enters the CURRENT state, for example after a data indication.
This can be fixed by allowing more time for the receive machine to receive the inbound packet.
CommentEnd:

SuggestedRemedy:
Increase the current_while_timer start value by a factor to reflect both

the asynchronous nature of the packets and the possible loss or delay of a PDU.
The timer should be started with the value 3 x Long_Timeout_Time or 3 x Short_Timeout_Time.

The transmit engine will still transmit at the faster rate.
RemedyEnd:
Editor'sRecommendationStart:

Reject. The values of timeout are already 3 times the values of the corresponding periodic TX timers. That is all that is needed.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 72
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.11
Page: 132
Line: 40
CommentType (E, T, ER, or TR): E
Comment:
Change "updated_selected" to "update_Selected"
CommentEnd:

SuggestedRemedy:
Change "updated_selected" to "update_Selected"
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 73
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408

CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.11
Page: 132
Line: 42
CommentType (E, T, ER, or TR): E
Comment:
Change "updated_matched" to "update_Matched"
CommentEnd:

SuggestedRemedy:
Change "updated_matched" to "update_Matched". (3 times: line 42, 44 and page 133 line 1).
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 74
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.4.11
Page: 133
Line: 11
CommentType (E, T, ER, or TR): TR
Comment:

The requirement not to validate the destination address is inappropriate - see my comment on 43.2.9; it also has nothing to do with future-proofing.

CommentEnd:

SuggestedRemedy:

Delete "Destination Address," from the list.

RemedyEnd:
Editor'sRecommendationStart:

Discuss. This wording was included specifically as a result of input from the last meeting and comment resolution.

Editor'sRecommendationEnd:
ResolutionStart:

Resolved by earlier comment.

ResolutionEnd:

CommentNumber: 75
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.11
Page: 133
Line: 13
CommentType (E, T, ER, or TR): T
Comment:
"A Receive Machine may validate the Length field".
Which "Length" field? (There are 4 Length fields).
CommentEnd:

SuggestedRemedy:
Specify which Length field may be verified.

It may help to give a unique name to each Length Field
(eg. Actor_Length, Partner_Length, Terminator_Length)?
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Add qualifiers to length fields; Receiver may validate any of the TLV length fields in the LACPDU

ResolutionEnd:

CommentNumber: 76
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.10
Page: 133
Line: 43
CommentType (E, T, ER, or TR): E

Comment:
spelling of Paerter rather than partner and wrong name anyway needs to be the port version
CommentEnd:

SuggestedRemedy:
Use Partner_Oper_Port_State.Aggregation
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 77
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.11
Page: p133
Line: 43
CommentType (E, T, ER, or TR): E
Comment:
Replace "Paertner_Oper_State.Aggregation" with Partner_Oper_Port_State.Aggregation".
CommentEnd:

SuggestedRemedy:
Replace "Paertner_Oper_State.Aggregation" with Partner_Oper_Port_State.Aggregation".
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 78
CommenterName: Mick Seaman

CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine
Page: 132
Line: 40
CommentType (E, T, ER, or TR): ER
Comment:
TA. The notion that the 'received protocol information is no longer valid' has been tighten up considerably in the last two drafts and is now concisely expressed as a change in the LAG. This is what is expressed in the update_Selected function. The vague words in this subclause are no longer appropriate. Similarly the choose_Matched function is better expressed in terms of agreeing a unique LAG. It is also unnecessary to include a summary of the detail of the choose_Matched function as the function is already described twice (page 129, lines 24-34) and page 131 lines 2-12). Note however that the first of those descriptions is inaccurate and should also be replaced by a shorter statement and a reference to the second, which mirrors the correct text in this subclause (43.4.11).
CommentEnd:

SuggestedRemedy:
Replace the following:
"The update_selected function is used to set the Selected variable to FALSE if the received protocol information indicates that the currently selected Aggregator is no longer valid. The choose_matched function determines whether the participants have both agreed upon the protocol information that they have exchanged, to the extent that a port can safely be used in an aggregate. The choose_matched function sets the Matched variable to TRUE if either of the following are true:

a) The Partner's Aggregation state is Individual (i.e., the value of Actor_State.Aggregation in the received LACPDU is FALSE); or
b) The Partner's view of the Actor's state, as represented in the Partner parameters in the received LACPDU, matches the Actor's operational parameters (Actor_Port_Number, Actor_Port_Priority, Actor_System_ID, Actor_System_Priority, Actor_Oper_Key, Actor_Oper_Port_State. Aggregation).

Otherwise, the choose_matched function sets the Matched variable to FALSE."

With the following:
"The update_Selected function sets the Selected variable FALSE if the Link Aggregation Group identified by the combination of the protocol partner's own information and the actor's own information has changed. If the same unique Link Aggregation Group is

correctly identified by the information in the LACPDU the choose_Matched function sets the Matched variable TRUE, and FALSE otherwise.

The Matched variable is used by the Mux Control and Logic (43.4.14).
RemedyEnd:
Editor'sRecommendationStart:

AIP. The editor should wordsmith the proposed text.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 79
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine
Page: 133
Line: 1
CommentType (E, T, ER, or TR): ER
Comment:
If the use of the matched variable is to be described here, so to be consistent, should the use of Selected.
CommentEnd:

SuggestedRemedy:
Before the text:
"The Matched variable is used by the Mux Control and Logic (<reference>)."
Add:
"The Selected variable is used by the Selection Logic and Machine (<reference>)."
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 80
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine
Page: 133
Line: 1
CommentType (E, T, ER, or TR): ER
Comment:
If the use of the matched variable is to be described here, so to be consistent, should the use of Selected.
CommentEnd:

SuggestedRemedy:
Before the text:
"The Matched variable is used by the Mux Control and Logic (<reference>)."
Add:
"The Selected variable is used by the Selection Logic and Machine (<reference>)."
RemedyEnd:
Editor'sRecommendationStart:

See comment 79.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 81
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine
Page: 133
Line: 10
CommentType (E, T, ER, or TR): TR
Comment:
The requirement not to validate the Destination Address is equivalent to a requirement to promiscuously receive and filter on Ethertype alone. This is an unnecessary and unrealistic burden to place on both end stations (most particularly) and bridges. It blows away the guarantees as to processing rate provided by the

Slow Protocols Multicast Address. It also is a significant revision to the purpose of having destination addresses at all in the 802 architecture. Remove it.
Additionally it appears from this paragraph that the Receive Machine will be compliant if it validates no fields at all, i.e. it should process every frame coming down the link. I am sure that was not what is intended. However the "may validate the Length field" is ambiguous. There is one Length/Type field and two Length fields in a LACPDU, how many of these are being referred to by "the Length field".

CommentEnd:

SuggestedRemedy:
Remove "destination Address," from line 11. This also requires accompanying changes in 43.2.9.
Clarify the reference to the Length field.
RemedyEnd:
Editor'sRecommendationStart:

See comments 74 and 75.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 82
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine
Page: 133
Line: 19
CommentType (E, T, ER, or TR): E
Comment:
In words "temporary holding state" give no clue as to what is being held. I don't think anything is.
CommentEnd:

SuggestedRemedy:
Replace "temporary holding state" by "transient state".
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 83
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.4.11
Page: 133
Line: 24
CommentType (E, T, ER, or TR): E
Comment:

The name used in this text for record_default_PDU do not match the name described in 43.4.8

CommentEnd:

SuggestedRemedy:

Make them the same in one place or the other.

RemedyEnd:
Editor'sRecommendationStart:

Accept. Use recordDefaultPDU consistently.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 84
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine
Page: 133
Line: 25
CommentType (E, T, ER, or TR): E
Comment:

TA. It is unclear what the utility of "these circumstances" is. It should be unnecessary to repeat the word "administrative". The phrase "verify whether the current selection state needs to be changed as a result of ..." can be made much more specific.

The name of the functions used should line up with the names used on page 130. Separately the names of these functions should be made consistent as to whether they contain underscores or not.
CommentEnd:

SuggestedRemedy:
Replace
"The current operational parameters for the Partner are overwritten by the record_default_PDU function, using the corresponding administrative values, allowing administrative configuration of aggregations and individual links under these circumstances. The update_default_selected function is used to verify whether the current selection state needs to be changed as a result of the default settings of the Partner variables. As the state machine is now set to a locally administered state, the Matched variable is set to TRUE."

With
"The recordDefaultPDU function overwrites the current operational parameters for the Partner with administratively configured values. This allows configuration of aggregations and individual links when no protocol partner is present, while still permitting an active partner to override default settings. The update_Default_Selected function sets the Selected variable FALSE if the Link Aggregation Group has changed. Since all operational parameters are now set to locally administered values there can be no disagreement as to the Link Aggregation Group, so the Matched variable is set TRUE."
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 85
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43

Subclause: 43.4.13
Page:
Line:
CommentType (E, T, ER, or TR): ER
Comment:
TA. This clause interleaves the description of the Selection Logic and the Selection Machine in a way that makes it hard for the reader to focus on what is going on. It opens with an injunction for the Selection Machine to implement the function specified in Figure 43-13, which does not appear for another four pages. The clause and its subclauses should be reorganized to:
a) Briefly identify what the Selection Logic and Selection Machine are.
b) Postpone the description of the states of the Selection Machine until subclause 43.4.13.3 which contains the figure.
c) Move as much detail about the considerations pertaining to the recommended default of the selection logic into the subclause dealing particularly with that subject (43.4.13.2) as possible.
d) Move detail from 43.4.13.2 that is applicable to all selection rules and is required for correct operation into 43.4.13.1.

The following trivial misstatements should also be removed or corrected:
e) The selection logic does not determine the LAG ID for the port. The LAG ID is already present in the data held by the Receive Machine.
f) The introductory description of the selection machine should not assume that a port is always attached to an "existing Aggregator", since it has just been explained that it may not be possible to attach a port to an Aggregator for various reasons.
g) It is the selection logic that determines that a port should be detached from an Aggregator, not the selection machine. The selection logic does not determine the correct Aggregator until the port has been detached from the prior Aggregator (other Aggregators may become free in this interval, other ports that have detached may have taken some).
h) A port must be detached from its Aggregator when the LAG ID changes, even if the same Aggregator is selected later (and, see above, the final Aggregator is not always known). To do otherwise is to risk misdelivery of frames. The condition for detaching from an Aggregator is therefore simply

CommentEnd:

SuggestedRemedy:
This suggested Remedy retains as much of the existing text as possible and is not intended to introduce any technical change.

Replace Clause 43.4.13 with the following:

"43.4.13 Selection Logic and Machine

The Selection Logic selects a compatible Aggregator for a port using the port's LAG ID. The selection logic may determine that the link should be operated as a standby if there are constraints on the simultaneous attachment of ports that have selected the same Aggregator. If the

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

The Selection Machine controls the process of attaching the port to a selected Aggregator, first detaching it from any prior Aggregator if the LAG ID has changed."

Sorry, I haven't had time to finish this.

RemedyEnd:
Editor'sRecommendationStart:

AIP. The editor should wordsmith the proposed text.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 86
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.13
Page: 134
Line: 52
CommentType (E, T, ER, or TR): E
Comment:
It might be good for clarity to add that the Selection variable is also modified by the receive machine in this paragraph.
CommentEnd:

SuggestedRemedy:
Add to the end of the paragraph on page 135 line 3:
The Selection variable is modified by the receive machine when the receive machine detects from received LACPDUs that the link is nolonger

selected.
RemedyEnd:
Editor'sRecommendationStart:

AIP. Reconcile the proposed text with the resolution of comment 85.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 87
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.13.1
Page: 135
Line: 48
CommentType (E, T, ER, or TR): E
Comment:
Partner ID should be Partner System ID
CommentEnd:

SuggestedRemedy:
Change Partner ID to Partner System ID
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 88
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.4.11
Page: 133
Line: 43
CommentType (E, T, ER, or TR): E

Comment:

Typo on Paertner_Oper_State

CommentEnd:

SuggestedRemedy:

Change to Partner_Oper_State

RemedyEnd:
Editor'sRecommendationStart:

Accept. See comment 76.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 89
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.13.1
Page: 135
Line: 47
CommentType (E, T, ER, or TR): T
Comment:
Add Actor (System?) ID before Actor Key.
CommentEnd:

SuggestedRemedy:
Add Actor (System?) ID before Actor Key.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 90

CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.4.13.3
Page: 137
Line: 25
CommentType (E, T, ER, or TR): E
Comment:

Spurious comma at the end of the line

CommentEnd:

SuggestedRemedy:

Remove the extra comma.

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 91
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.13.3
Page: 137
Line: 26
CommentType (E, T, ER, or TR): T
Comment:
Should "Status.aggregate" be in the list of things that
may change too?

Also, spurious comma at end of line 25.
CommentEnd:

SuggestedRemedy:
Add "Status.aggregate" to list of things that may change.

Remove spurious comma at end of line 25.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 92
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.13.3
Page: 137
Line: 30
CommentType (E, T, ER, or TR): E
Comment:
Remove extra "selecting".
CommentEnd:

SuggestedRemedy:
Remove extra "selecting".
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 93
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.13.3
Page: 137
Line: 37
CommentType (E, T, ER, or TR): E

Comment:
List begins with d).
CommentEnd:

SuggestedRemedy:
Start new list.
RemedyEnd:
Editor'sRecommendationStart:

Reject. See comment 8.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 94
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.13.3
Page: 138
Line: 10
CommentType (E, T, ER, or TR): E
Comment:
Replace "DTACHING" with "DETACHING".
CommentEnd:

SuggestedRemedy:
Replace "DTACHING" with "DETACHING".
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 95
CommenterName: Graham Short
CommenterEmail: Graham.Short@tiuk.ti.com
CommenterPhone: +44-1604-663410
CommenterCompany: Texas Instruments

Clause: 43
Subclause: 43.4.13.3
Page: 138
Line: 10
CommentType (E, T, ER, or TR): E
Comment:
typo
CommentEnd:

SuggestedRemedy:
Change DTACHING to DETACHING
RemedyEnd:
Editor'sRecommendationStart:

Accept. See comment 94.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 96
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.13.3
Page: 138
Line: 10
CommentType (E, T, ER, or TR): E
Comment:
DTACHING should be DETACHING
CommentEnd:

SuggestedRemedy:
replace DTACHING with DETACHING
RemedyEnd:
Editor'sRecommendationStart:

Accept. See comment 94.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 97
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.4.13.3
Page: 138
Line: 10
CommentType (E, T, ER, or TR): E
Comment:

Typo on DTACHING

CommentEnd:

SuggestedRemedy:

Insert in E in-between D and T. :-)

RemedyEnd:
Editor'sRecommendationStart:

Accept. See comment 94.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 98
CommenterName: Peter Saunderson
CommenterEmail: Peter_Saunderson@3Com.com
CommenterPhone: +44 1442 438084
CommenterCompany: 3Com
Clause: 43
Subclause: 43.4.14
Page: 139
Line: 8
CommentType (E, T, ER, or TR): T
Comment:
Point a) and b) Not a full description of the IN_SYNC / OUT_OF_SYNC states - too wordy and open to misinterpretation. For example it does not state that when the Matched variable is FALSE then the Mux is OUT_OF_SYNC. It is not clear on the value of the Partners state needed for the Mux to be IN_SYNC.
CommentEnd:

SuggestedRemedy:

Replace the bullets and Note 2 with the following:

a) The Mux is IN_SYNC if the Selection Machine is in the ATTACHED state having attached the port to the correct Aggregator, and the Receive machine indicates Matched = TRUE. The Mux is OUT_OF_SYNC if the conditions for IN_SYNC are not met. The actor must signal In_Sync or Out_Of_Sync in the transmitted LACPDUs depending on the state of the Mux.
b) If the Mux is OUT_OF_SYNC or the Partner's state is signaling Out_Of_Sync, then both the collector and distributor should be turned off.

RemedyEnd:
Editor'sRecommendationStart:

AIP. The editor should wordsmith the proposed text.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 99
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.4.14
Page: 139
Line: 25
CommentType (E, T, ER, or TR): E
Comment:
List begins with g).
CommentEnd:

SuggestedRemedy:
Start new list.
RemedyEnd:
Editor'sRecommendationStart:

Reject. See comment 8.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 100
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.5.1
Page: 142
Line: 6
CommentType (E, T, ER, or TR): E
Comment:

Excessive use of semicolons.

CommentEnd:

SuggestedRemedy:

Change "...protocol is optional; some distribution..." to "...protocol is optional. Some distribution..."

RemedyEnd:
Editor'sRecommendationStart:

Accept;

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 101
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.5.3.2
Page: 143
Line: 15,17,19,46,48,49
CommentType (E, T, ER, or TR): E
Comment:
Replace "Requestor" with "Requester".
CommentEnd:

SuggestedRemedy:
Replace "Requestor" with "Requester". (9x)
RemedyEnd:

Editor'sRecommendationStart:

Reject unless the TF insists on acceptance. Both forms are considered acceptable.

Editor'sRecommendationEnd:
ResolutionStart:

Accept. Global search and replace.

ResolutionEnd:

CommentNumber: 102
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.5.3.2
Page: 143
Line: 32
CommentType (E, T, ER, or TR): ER
Comment:

The use of Slow_Protocols_Multicast address in bullet a) does not align with terminology in 43B.

CommentEnd:

SuggestedRemedy:

Replace "Slow_Protocols_Multicast address" with "Slow Protocols group MAC Address".

RemedyEnd:
Editor'sRecommendationStart:

Align with resolution to comment 55.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 103
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.5.3.2

Page: 143
Line: 35/36
CommentType (E, T, ER, or TR): ER
Comment:

The use of Slow_Protocol type field in bullet c) does not align with terminology in 43B.

CommentEnd:

SuggestedRemedy:

Replace "Slow_Protocol type field" with "Slow Protocols Ethernet Type"

RemedyEnd:
Editor'sRecommendationStart:

Align with resolution of comment 56.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 104
CommenterName: Paul Congdon
CommenterEmail: paul_congdon@hp.com
CommenterPhone: (916) 785-5753
CommenterCompany: Hewlett Packard
Clause: 43
Subclause: 43.5.3.1
Page: 144
Line: 3
CommentType (E, T, ER, or TR): E
Comment:

A description of how the reserved field is to be treated should be included as it was with the reserved bits in the State field of LACPDU frames.

CommentEnd:

SuggestedRemedy:

Indicate the padded fields are to be echoed as received, or some other action to take when reflecting the frame. If it is not necessary to echo the pads, then the instructions in the RESPOND state of figure 43-19 are too specific.

RemedyEnd:
Editor'sRecommendationStart:

Discuss.

Editor'sRecommendationEnd:
ResolutionStart:

Clarify: Marker Requests must transmit reserved field as zeroes. Add wording to specifically allow an implementation to reflect the received value or transmit zeroes in a Marker Response.

ResolutionEnd:

CommentNumber: 105
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: 43.5.4.1
Page: 144
Line: 18
CommentType (E, T, ER, or TR): E
Comment:

The subclause is describing protocol operation.

CommentEnd:

SuggestedRemedy:

Change the title from "Operation of the marker service" to "Operation of the marker protocol".

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 106
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd

Clause: 43
Subclause: 43.5.4.1
Page: 144
Line: 25
CommentType (E, T, ER, or TR): E
Comment:
Replace "Requestor" with "Requester".
CommentEnd:

SuggestedRemedy:
Replace "Requestor" with "Requester". (3x)
RemedyEnd:
Editor'sRecommendationStart:

Align with resolution of comment 101.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 107
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.5.4.1
Page: 144
Line: 32
CommentType (E, T, ER, or TR): TR
Comment:
The 5 frames/second limit needs to be clarified. If
an Actor is issuing 5 Marker frames/second then the Partner
needs to send 5 Response frames/second. If it is doing this
then how can it also send 5 Marker frames of its own since it is
already at the Marker Protocol's limit?

Need to allow 5 Markers and 5 Reponses to be issued per second.
This then however breaks the 5 frames/second total per Slow Protocol.
CommentEnd:

SuggestedRemedy:
Possibilities....
* Raise limit per slow protocol to 10. Limit to 5 each for
Marker and Response for this standard.
* Redefine Marker and Response to be different Slow Protocols.
* Remove the Marker/Response pototocol by exchanging collector
delays via the LACP.

RemedyEnd:
Editor'sRecommendationStart:

Discuss. The second option appears to be the easiest, however it implies changing the encoding of the Marker vs. the Marker Response PDUs (different slow protocols take different protocol subtype values). Also, it would consume 3 of the 10 slow protocols provided.

Editor'sRecommendationEnd:
ResolutionStart:

AIP. Clarify wording such that Marker Response don’t count towards the SlowProtocol Limit--only Markers do. Add text at p144 l32 that, if an implementation IS limiting its responses, that the superfluous Markers are dropped and the responses are not indefinitely queued.

Consider reopening this one; a possible fix is to modify the SlowProtocols definition itself.

ResolutionEnd:

CommentNumber: 108
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43
Subclause: Fig 43-19
Page: 145
Line: 24/25
CommentType (E, T, ER, or TR): TR
Comment:

The Note attached to this diagram expresses a mandatory requirement.

CommentEnd:

SuggestedRemedy:

Elevate the Note to the status of normative text.

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 109

CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.7.4
Page: 149
Line: 20
CommentType (E, T, ER, or TR): E
Comment:
There are two FC1s.
CommentEnd:

SuggestedRemedy:
Renumber.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 110
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.7.4
Page: 149
Line: 20
CommentType (E, T, ER, or TR): E
Comment:
Wrong subclause reference. Should be 43.2.3.1.4
CommentEnd:

SuggestedRemedy:
Change reference to 43.2.3.1.4
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 111
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.7.9
Page: 151
Line: 6
CommentType (E, T, ER, or TR): E
Comment:
The "shall" in 43.2.10 p110 line 36 does not seem to have a PICS entry except that the sentence is virtually the same as the one in 43.3.3 p112 line 27, which does have a PICS entry.
CommentEnd:

SuggestedRemedy:
Add another PICS entry, or clean up text to only have one description (or do nothing if considered okay).
RemedyEnd:
Editor'sRecommendationStart:

Add a PICS entry.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 112
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.3.7.10(b?)
Page: 151
Line: 27
CommentType (E, T, ER, or TR): E
Comment:
The "shall" in 43.3.4 on p112 line 45 does not appear to have a PICS entry.
CommentEnd:

SuggestedRemedy:

Add PICS entry.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 113
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.7.11
Page: 151
Line: 34
CommentType (E, T, ER, or TR): E
Comment:
Wrong subclause reference. Should be 43.3.5.
CommentEnd:

SuggestedRemedy:
Change 43.3.4 to 43.3.5.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 114
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.7.15
Page: 152
Line: 45
CommentType (E, T, ER, or TR): E

Comment:
There is a "shall not" in 43.4.11 on p133 line 10 which has no corresponding PICS entry.
CommentEnd:

SuggestedRemedy:
Add PICS entry.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 115
CommenterName: Graham Short
CommenterEmail: Graham.Short@tiuk.ti.com
CommenterPhone: +44-1604-663410
CommenterCompany: Texas Instruments
Clause: 43
Subclause: 43.7.17
Page: 153
Line: 16
CommentType (E, T, ER, or TR): E
Comment:
Rerence to incorrect Figure
CommentEnd:

SuggestedRemedy:
Change Figure 43-12 to Figure 43-13. The figure quoted is the recommended default which is not mandatory as SLM1 is.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 116
CommenterName: Keith.Balmer

CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.7.18
Page: 154
Line: 6
CommentType (E, T, ER, or TR): E
Comment:
Wrong font used for "Mux Control and Logic".
CommentEnd:

SuggestedRemedy:
Change font.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 117
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 154
Page: 43.7.19
Line: 40
CommentType (E, T, ER, or TR): E
Comment:
The "shall" in 43.4.15 on p140 line 3 requiring NTT
to be set false does not have an associated PICS entry.
CommentEnd:

SuggestedRemedy:
Add PICS entry.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 118
CommenterName: Graham Short
CommenterEmail: Graham.Short@tiuk.ti.com
CommenterPhone: +44-1604-663410
CommenterCompany: Texas Instruments
Clause: 43
Subclause: 43.7.21
Page: 155
Line: 8
CommentType (E, T, ER, or TR): E
Comment:
Reference to 43.5.2 but I couldn't find any mention here (or in other sections for that matter) about the maximum request rate of five per second.
CommentEnd:

SuggestedRemedy:
Either remove or correct reference or add data to 43.5.2 for the request rate.
RemedyEnd:
Editor'sRecommendationStart:

Change reference to 43.5.4.1

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 119
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43.7.21
Page: 155
Line: 8
CommentType (E, T, ER, or TR): E
Comment:
Wrong subclause reference (43.5.2). Possibly meant to be 43.5.4.1?
CommentEnd:

SuggestedRemedy:
Change to correct reference (43.5.4.1?).
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 120
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43A
Subclause: 43A.3
Page: 159
Line: 18
CommentType (E, T, ER, or TR): E
Comment:
Change "describes" to "described".
CommentEnd:

SuggestedRemedy:
Change "describes" to "described".
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 121
CommenterName: Tony Jeffree
CommenterEmail: tony@jeffree.co.uk
CommenterPhone: +44-161-973-4278
CommenterCompany: Independent Consultant
Clause: 43B
Subclause: 43B.4
Page: 163

Line: 38, 43
CommentType (E, T, ER, or TR): E
Comment:

Spurious quotation marks in Table 43B-2(consistency with Table 43B-1).

CommentEnd:

SuggestedRemedy:

Lose the quotes around "Slow Protocols" in the table title and in the body of the table. Check for other occurrences.

RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 122
CommenterName: Keith.Balmer
CommenterEmail: Keith.Balmer@tiuk.ti.com
CommenterPhone: (44) 1604 663408
CommenterCompany: Texas Instruments Ltd
Clause: 43
Subclause: 43B.6.2.?
Page: 165
Line: 51
CommentType (E, T, ER, or TR): E
Comment:
Where are the PICS entries for 43B? There are 5 "shalls" but no PICS entries.
CommentEnd:

SuggestedRemedy:
Add the PICS entries.
RemedyEnd:
Editor'sRecommendationStart:

Accept.

Editor'sRecommendationEnd:
ResolutionStart:

Follow Editor's Recommendation

ResolutionEnd:

CommentNumber: 123
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.8 and Figure 43-10 state diagram
Page: 129
Line: 44-54
CommentType (E, T, ER, or TR): TR
Comment:
The function update_LACPDisabled_Selected goes to some trouble to effect a graceful transition to the LACP_DISABLED state, however returning to the LACP operational states is by way of INITIALIZE and always sets Selected False. There doesn?t seem much point to take great care in one direction of this rare transition and not in the other.
CommentEnd:

SuggestedRemedy:
In the interests of simplicity, get rid of the update_LACPDisabled_Selected function and replace it with Selected = False in the state diagram.
RemedyEnd:
ResolutionStart:

Accept

ResolutionEnd:

CommentNumber: 124
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.8
Page:
Line:
CommentType (E, T, ER, or TR): TR
Comment:
A better function for update_Selected is the following, which in combination with a better version of choose_Matched (see below), allows the link to be used with fewer messages exchanged in the case where on or both of the communicating systems have identified the link as an individual link (perhaps because the key is unique). In addition to this improvement this function sets the actor?s ?In_sync? state false along with Selected to prevent the

update_NTT function from incorrectly asserting that the partner has the actor?s ?In sync? wrong. Of course the Mux control logic would / will also do this, but there remains the question of the relative speed of the infinitely fast machines as to whether Mux would get there in time.

The choose_Matched function also recognizes that the current match will (highly probably) shortly come to an end if both actor and partner are Passive, and hence holds off the match. Once both Passive partners time out each other it is possible that the default (admin) parameters will lead to an aggregation.

```
static void update_selected(Lac_port *port, Lac_pdu *pdu)
{
  if ( (   same_partner(&pdu->actor, &port->partner)
       &&  pdu->actor.state.aggregation
       &&  port->partner.state.aggregation
       &&  port->actor.state.aggregation
       )
    | | ( !pdu->actor.state.aggregation
       &&(!port->partner.state.aggregation | | !port-
>actor.state.aggregation)
       )
    | | ( !port->actor.state.aggregation
       && !pdu->partner.state.aggregation
       &&(!port->partner.state.aggregation | | !port-
>actor.state.aggregation)
    ) )
  {}
  else
  {
    port->selected = False;
     port->actor.state.synchronization = False;
} }

static void choose_matched(Lac_port *port, Lac_pdu *pdu)
{
  if ( ( (   same_partner(&pdu->partner, &port->actor)
         && (pdu->partner.state.aggregation ==
            port->actor.state.aggregation)
        )
       | | ( !pdu->actor.state.aggregation
        )
       | | ( !port->actor.state.aggregation
         && !pdu->partner.state.aggregation
      ) )
    && ( (   pdu->actor.state.lacp_activity
        )
       | | (   port->actor.state.lacp_activity
          &&  pdu->partner.state.lacp_activity
    ) ) )
```

```
    port->matched = True;
  else
    port->matched = False;
}
```

CommentEnd:

SuggestedRemedy:
Upgrade the update_selected and chose_Matched functions as suggested, converting the code into appropriately unambiguous English.
RemedyEnd:
ResolutionStart:

AIP. Mick to update submission and provide appropriate text.

ResolutionEnd:

CommentNumber: 125
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.8 Receive Machine
Page: 131
Line: 20
CommentType (E, T, ER, or TR): TR
Comment:
The update_NTT function should also catches differences in LACP Activity, LACP Timeout, Aggregation, Synchronization (with appropriate care as commented on further below), and Collecting.
CommentEnd:

SuggestedRemedy:
Add the above flags for checking in this function.
RemedyEnd:
ResolutionStart:

Accept. Add text to p131 l14-22

ResolutionEnd:

CommentNumber: 126
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine, Figure 43-10 state diagram

Page:
Line:
CommentType (E, T, ER, or TR): TR
Comment:
The port_moved variable, if asserted will cause the machine to continuously circulate through the INITIALIZE and PORT_DISABLED states. This is unsatisfactory for comprehension, and also raises the question as to what happens if both port_moved and port_enabled are TRUE. As a further minor point the state of the port_moved variable is not declared.
CommentEnd:

SuggestedRemedy:
Replace ?port_moved? by ?port_moved = True? and set ?port_moved = FALSE? in INITIALIZE.
RemedyEnd:
ResolutionStart:

Withdrawn.

ResolutionEnd:

CommentNumber: 127
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine, Figure 43-10 state diagram
Page:
Line:
CommentType (E, T, ER, or TR): TR
Comment:
It is unclear what is meant to happen if both LACP_Enabled = FALSE and Reinitialize = TRUE.
CommentEnd:

SuggestedRemedy:
AND the current global transition for LACP_Enabled = False with Reinitialize = FALSE, so it is clear which of these takse precedence. Qualify the port_enabled = TRUE transition from PORT_DISABLED by ANDing it with LACP_Enabled = TRUE.
RemedyEnd:
ResolutionStart:

AIP. Qualify Reinitialize=TRUE with LACP Enabled=TRUE

ResolutionEnd:

CommentNumber: 128
CommenterName: Mick Seaman

CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine, Figure 43-10 state diagram
Page:
Line:
CommentType (E, T, ER, or TR): TR
Comment:
A separate LACP_DISABLED state (as a resting state)is not required. What is more, transition to this state by setting LACP_Enabled = False risks misdelivery of frames. A better solution to the requirement to turn LACP off, even if the other end of the link is trying to talk LACP, is to use a lacp_enabled = FALSE condition to discard all received frames while at the same time setting Matched false and setting both actor and partner administrative and operational states to Passive and Individual. That should cause the receive machine to transition to the DEFAULTED state with sufficient delay for the other end of the link to do whatever it does when it has a non-LACP capable partner.
CommentEnd:

SuggestedRemedy:
Remove the LACP_DISABLED state (or just use it to set conditions with a UCT to CURRENT) but retain the LACP_Enabled variable. Use this variable as a condition on the receive LACPDU events. Add management constraints on the values of other parameters if LACP Enabled is ever set FALSE to produce the desired least damaging behavior in the DEFAULTED state.
RemedyEnd:
ResolutionStart:

Withdrawn

ResolutionEnd:

CommentNumber: 129
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.11 Receive Machine, Figure 43-10 state diagram
Page: 132
Line: 11
CommentType (E, T, ER, or TR): TR
Comment:
CommentEnd:
Setting Actor_Oper_Port_State.LACP_Timeout = Short_timeout in EXPIRED makes the protocol significantly more chatty (~50%) on

startup if NTT is asserted when the partner?s view of LACP_Timeout is incorrect. I believe the latter is desirable. Message loss on startup to an actor running long timeouts is already protected since the actor will retry (Partner_Oper_Port_State.LACP_Timeout = Short_timeout) and provoke another response from the partner by claiming an incorrect partner id.

SuggestedRemedy:
Remove the change to the Actor?s timeout in EXPIRED. Keep the change to the Partner?s timeout.
RemedyEnd:
ResolutionStart:

Accept. Also remove the Actor_Oper_Port_State assignment from the CURRENT state.

ResolutionEnd:

CommentNumber: 130
CommenterName: Mick Seaman
CommenterEmail: mick_seaman@ieee.org
CommenterPhone: 650 796 4836
CommenterCompany: 3Com Corporation
Clause: 43
Subclause: 43.4.15
Page: 139
Line: 26
CommentType (E, T, ER, or TR): ER
Comment:
NTT should not be asserted on an IN_SYNC to OUT_OF_SYNC transition if the cause of that transition is Selected becoming false on receipt of a PDU from a partner. In that case the partner already knows that any prior received true value of ?Synchronization? is invalid.
CommentEnd:

SuggestedRemedy:
Handle the above case in the receive machine so that the mux doesn?t see it.
RemedyEnd:
ResolutionStart:

Withdrawn.

ResolutionEnd:

CommentNumber: 131
CommenterName: Hadriel Kaplan
CommenterEmail: hsk@iol.unh.edu
CommenterPhone: (603)862-3332
CommenterCompany: UNH InterOperability Lab

Clause: 43
Subclause: 43.4.4
Page: 123
Line: 52
CommentType (E, T, ER, or TR): TR
Comment:
The timer constants have no max aand min values. I realize they're implemented in software and don't need to be very tight, but they do need to be bound. There are 2 reasons for this:
1) If an implementation has a slow fast_periodic timer and only sends LACPDUs once every 1.5 seconds, and another implementation has a fast short_timeout timer and invalidates PDUs after only 2.5 seconds, then even one dropped LACPDU will cause the timeout to occur. Without a max and min bound, implementers can use essentially whatever values they want.
2) The second reason is political: I believe thee rest of 802.3 will want timer boundaries and so this will just be addressed later - might as well do it now.
CommentEnd:

SuggestedRemedy:
Either specify particular max and min values for each timer, or globally specify a +/- 10% range.
RemedyEnd:
ResolutionStart:

AIP. Add a tolerance of +/- 250 ms to all timers here.

ResolutionEnd:

CommentNumber: 132
CommenterName: Hadriel Kaplan
CommenterEmail: hsk@iol.unh.edu
CommenterPhone: (603)862-3332
CommenterCompany: UNH InterOperability Lab
Clause: 43
Subclause: 43.2.4.1.1
Page: 105
Line: 41-44
CommentType (E, T, ER, or TR): ER
Comment:
Throughout the draft the variables "DA", "SA", "m_sdu", and "service class" are listed without definitions.
CommentEnd:

SuggestedRemedy:
Add their respective definitions in all locations (do global search).
RemedyEnd:
ResolutionStart:

Reject in principle. These are all defined in the MAC Service Definition in Clause 2. Add the phrase “as defined in Clause 2” to the definition.

ResolutionEnd: