IEEE 802 Plenary Meeting

Austin, TX

8th - 12th March '98

1. Discussion Details for 802.3ad

1.1 Intro

Monday 8th March, 8.00. Meeting was called to order at 8:15.

This morning meeting is treated as an interim, as it is being held prior to the opening plenary.

About 400 comments received. Steve Haddock presented a summary of these by section. This indicated that the biggest issues were to do with key changes and aggregate identifiers. It was decide to address these issues first.

1.2 Combining Some State Machines - Mick Seaman

Mick has proposed that there are two main components: the front end protocol handling; and the back end aggregator control. Also includes the definition of a complete set of default information for the protocol partner, such that only the receive handler needs to know where this information came from. This simplifies the amount of checking the rest of the system needs to do. This allows some of the state machines to be combined and simplified.

Steve Haddock suggested we need a definition of the combined state machines, and a list of the comments that are addressed by this. Mick Seaman, Tony Jeffree, Peter Saunderson, Hadriel Kaplan, Benjamin Brown, Joris Wils and Stephen Haddock to work offline to produce this.

1.3 Comments

Comment 166: Mick Seaman

Mick does not like reference to 'LAN Aggregation' which is not defined here. He wants specific exclusion of shared media, even though this is implied by a) and c), as this is the real intention. Suggestion is to make the bullet title "a) Aggregation across groups of LANs" with the current body text unchanged.

Comment 168: Mick Seaman

Accepted.

Comment 167: Hadriel Kaplan

1) Suggestion is to line up Control and Aggregator Parsers and show a point to point connection for at least one of the pairs with a not to indicate analgous connections for the others. Eliminate the BUS connections.
2) Show two lines between Control and Aggregator blocks.
3) Accept.

Comment 173: Robin Tasker

Reject. Change "will bind" to "binds".

Comment 174: Robin Tasker

Accept.

Comment 175: Robin Tasker

Accept in principle. Add a bullet to show that a conversation can be switched to another link to maintain availability.

Comment 176: Robin Tasker

Add a reference.

Comment 177: Paul Congdon

Add “(for example Keys)” to text.

Comment 179: Paul Congdon

Accept in principle. Scratch the first occurrence of the word ‘unchanged’ from the second paragraph of 43.2.2. Add a sentence explaining the MAC address of the Aggregator will be inserted, if not provided by the client.

Comment 192: Hadriel Kaplan

Reject.

Comment 193: Hadriel Kaplan, Bob Noseworthy

Accept. Name each interface. Qualify each primitive usage with internal layer prefix. Editor to deal with this offline. Global change to Clause 43.

Comment 201: Norman Finn

Accept. Delete the note; capture and expand the text of the note in section 43.2.11.9.

Comment 207: Robin Tasker

Accept in principle. Delete the note.

Comment 211: Robin Tasker

Reject.

Comment 212: Devendra Tripathi

Reject.

Comment 208: Devendra Tripathi

Reject.

Comment 213: Norman Finn

Accept. Make it consistent that Aggregator ID is ‘ifIndex’ and Aggregator Address is specifically stated where the MAC address is required.

Action: Steve Haddock to allocate someone to this task.

Comment 214: Devendra Tripathi

Resolved by 213.

Comment 215: Stephen Haddock

There was some discussion as to whether we need to dynamically change keys, e.g. because Auto-Negotiation selects a different link speed. Consensus was YES, but we change the ‘operational’ key (not the ‘administrative’ key). There are no rules as to how this is achieved.

Accepted. The editors are to propose new text to introduce this section, with reference to 43.2.5 that defines how changes may occur.

Comment 217: Robin Tasker

Accept. Change ‘are used to determine’ to ‘may affect’ and give an example.

Comment 218: Robin Tasker

Accept. Resolved by 215.

Comment 216: Keith Balmer

Same as 218.

Comment 221: Hadriel Kaplan

Accept. Remove 2) line 8.

Comment 222: Hadriel Kaplan

Accept. In non-individual links, SK is sufficient, for Individual links SKPQ is required.

Comment 231: Hadriel Kaplan

Accept in principle. Remove ‘shall be flushed’; change to ‘frames in transit have been received by the collector’s receive function’. Rename the ‘Flush’ protocol as ‘Marker’ protocol. Let the editors explain ‘flush’ the first time it is used.

Comment 232: Keith Balmer

A problem with ‘flushing’ is the ordering constraints placed on the distributor take no account of the behaviour of the collector. Some discussion around limiting the time a collector can wait before treating a ‘received’ frame as ‘collected’. State diagrams mandate this is instantaneous.

Accept. Include a note that practical implementations may include queuing and delay in the Collector. Impose a time constraint on the Collector, that it must either deliver to its client, or discard frames within a CollectorMaxDelay (1 second). Note to account for propagation delay as well. Collectors can assume that frames have been received after a CollectorMaxDelay plus the propagation delay.

Comment 237: Stephen Haddock

Accept. Re-order fields within a TLV-tuple to align multi-byte fields on a 2-byte boundary.

Comment 238: Hadriel Kaplan

Accept. Ignore TL of TLV even if malformed. Receiver should be tolerant for these fixed version 1 fields.

Comment 239: Keith Klamm

Reject. Sender sets length=0, receivers can ignore length.

Comment 243: Tony Jeffrees

Accept.

Comment 244: Mick Seaman

Resolved by 243.

Comment 248, 251, 253: Mick Seaman

Deferred.

Comment 246: Joris Wils

Reject. The state machine runs continuously.

Comment 259: Mick Seaman

Accept. Eliminate LAG ID as a state variable.

Comment 260: Graham Short

Resolved by 259. Also eliminate Port_LAG_ID.

Comment 261: Stephen Haddock

Deferred.

Comment 263: Mick Seaman

Accept.

Comment 262: Stephen Haddock

Resolved by 263.

Comment 268: Mick Seaman

Resolved by 259.

Comment 270: Stephen Haddock

Resolved by ?.

Comment 273: Mick Seaman

Move Collector and Distributor state variables to the aggregator section.

Comment 275: Mick Seaman

Resolved by 259.

Comment 277: Mick Seaman

Accept. Change ‘Key’ to ‘OperKey’ for both Port_Key and Agg_Key. Text to explain how Admin Key maps to Oper Key.

Comment 278: Mick Seaman

Accept.

Comment 279: Mick Seaman

Accept. Add the variables and have the state machines show how the default is used or overridden by protocol exchanges.

Comment 283: Mick Seaman

Deferred.

Comment 285: Joris Wils

Will be resolved by 283.

Comment 287: Stephen Haddock

Will be resolved by 283.

Comment 289: Mick Seaman

Deferred.

Comment 290: Benjamin Brown

Accept.

Comment 294: Jeff Lynch

Accept. Global check for ‘transmission’ slow or fast, ‘receiver timeout’ as long or short.

Comment 295, 297: Mick Seaman

Deferred.

Comment 388: Devendra Tripathi

Reject. This is normal state machine convention.

Comment 390: Devendra Tripathi

Deferred.

Comment 302, 303, 304, 305, 306, 307, 308, 309, 310, 316, 318, 319, 320, 322, 323, 324, 326: …

Deferred.

Comment 327: Hadriel Kaplan

Reject.

Comment 329: Hadriel Kaplan

Deferred.

Comment 330: Hadriel Kaplan

Resolved by 215.

Comment 332: Norman Finn

Need both MAC address and ID versions of the Aggregator defined here. See also 213.

Comment 331: Keith Balmer

Accept. Make all bullets conformance requirements.

Comment 391: Hadriel Kaplan

Accept in principle. Resolution to be defined.

Comment 333: Hadriel Kaplan

Reject.

Comment 338: Joris Wils

Deferred.

Comment 339: Hadriel Kaplan

Accept.

Comment 340: Hadriel Kaplan

Accept in principle. Change the wait_while_timer to 2 seconds.

Comment 342: Joris Wils

Resolved by ?.

Comment 344: Hadriel Kaplan

Deferred.

Comment 345: Christian Thrysoe

Accept in principle. Add a sentence to e) that says when coupled, turn on Distributor and Collector when you are IN_SYNC.

Comment 347: Tony Jeffree

Accept.

Comment 348: Hadriel Kaplan

Accept. Merge note with normative text.

Comment 351: Keith Balmer

Accept. Add the statement to the end of this sub-clause.

Meeting adjourned at 1630 on Tuesday 9th March.

Meeting convened at 0900 on Wednesday 10th March.

Comment 393: Hadriel Kaplan

Accept in principle. Change implementation of Churn Machine sub-clause to be optional based on implementation of the related Management function. Fix text to say that this means that a port is up but has not been attached to an Aggregate Port within a bounded time. Change the timer to 60 seconds.

Comment 353: Hadriel Kaplan

Resolved by 393.

Comment 354: Keith Balmer

Accept.

Comment 394: Hadriel Kaplan

Accept in principle. Remove the unconditional entry of INFO_EXPIRED, add a condition on transfer to NO_REMOTE_CHURN on current_while_timer expired. (Problem with a race condition, take offline.) Subsequent resolution is to make Remote Churn Machine identical to Local Churn Machine.

Comment 356: Norman Finn

Accept. Reduce Pad to 2 bytes.

Comment 396: Hadriel Kaplan

Reject. (Make sure there is no validation performed on these fields. Define specifically what does need to be validated, just what has been previously done in the Parser, i.e. the Type and Subtype, no other fields are validated.)

General principle to apply to both Marker and LACPDUs: Specify in the clause where a frame is received which fields SHALL be verified, which SHALL NOT be verified, and which MAY be verified.

Comment 358: Hadriel Kaplan

Accept. Change to a 2-state diagram, indicate entry into a Wait_For state.

Comment 359: Hadriel Kaplan

Accept in principle. Add text to say that a Marker Responder MAY, but is not required to, throttle transmissions per the slow protocol timing constraint when faced with a Marker that does not comply.

Comment 362, 364: Norman Finn

Deferred. Resolve offline, Mick Seaman and Norm.

Comment 365: Hadriel Kaplan

Accept in principle. Modify Periodic state machine to provide a mechanism to prevent any LACPDU transmissions on a link. This mechanism will be used on shared media links. Assess scope of text changes for shared media accommodation, to indicate how the protocol gets "turned off".

Comment 6: Hadriel Kaplan

Accept.

Comment 7: Devendra Tripathi

Reject.

Comment 8: Hadriel Kaplan

Reject.

Comment 9: Tony Jeffree

Accept. Editors will do checking.

Comment 12: Tony Jeffree

Accept.

Comment 16, 17, 18: Les Bell

Reject.

Comment 19: Mick Seaman

Accept in principle.

Comment 20: Mick Seaman

Withdrawn.

Comment 21: Mick Seaman

Deferred.

Comment 23: Mick Seaman

Accept in principle.

Comment 29: Mick Seaman

?

Comment 35: Mick Seaman

Accept.

Comment 36: Paul Congdon

Reject. Port defaults are sufficient.

Comment 40: Norman Finn

Accept. Add the missing attribute.

Comment 44: Keith Klamm

Reject.

Comment 45: Keith Balmer

Accept.

Comment 46: Graham Short

See 45.

Comment 47: Keith Klamm

Accept. Make the text consistent with 30.3.1.1.8

Comment 50: Keith Klamm

Reject. Add a note to indicate why half-duplex may be auto-negotiated. An Aggregator can consist of a single half-duplex link, in which case a frame may discarded due to excessive collisions.

The architectural model, i.e. that an Aggregated Port looks like a conventional physical port, will be explained/ expanded on in 30.7.1.

1a. Agg layer doesn't look like a MAC from mgmt point of view.
Counts frames passing through.
Does not count any underlying errors or frames that bypass Agg layer.

1b. As 1a, add 'generic' summation of errors from below.

2. Agg layer looks like a MAC.
Count frames passing through top MAC Client i/f.
Count errors that occur in underlying MACs as if they occurred in Agg "MAC".

3. Approximate #2 by just summing (sum of deltas) the underlying MAC counters.
Make no attempt to subtract out LAC PDUs or frames bypassing Agg layer.

Straw poll: 1a 3; 1b 12, 2 18, 3 16.

Run off poll: 2 13, 3 11.

Stick with 2 (as we had previously defined). Note: do not differentiate between errors that occur on LAC PDUs or user data, but do differentiate for data counts. Explain this in 30.7.1.

Comment 57: Paul Congdon

Reject. The method of terminating the list is encoding dependent, but it exists. e.g. SNMP elements have a definite length.

Comment 59: Mick Seaman

Accept.

Comment 61: Mick Seaman

Accept.

Comment 65: Paul Congdon

Accept.

Comment 75: Mick Seaman

Accept.

Comment 77: Mick Seaman

Accept.

Comment 78: Mick Seaman

Accept.

Comment 79: Mick Seaman

Accept.

Comment 81: Tony Jeffree

Accept.

Comment 105: Tony Jeffree

Accept.

Comment 106: Tony Jeffree

See 105.

Comment 111: Les Bell

Accept b).

Comment 116: Jeff Lynch

Same as 111.

Comment 123: Jeff Lynch

Accept.

Comment 148: Les Bell

Accept.

Comment 156: Les Bell

Accept. b)

Comment 1: Keith Klamm

Add 'Actor' and 'Partner' to definitions.

Comment 2: Devendra Tripathi

Accept 2) literally.

Comment 3: Devendra Tripathi

Accept 2) - literally.

Comment 42: Mick Seaman

Accept in principle. Delete 'operational' and 'disabled'. Add text to Clause 43 explaining the behaviour, and add text here referring to this.

Meeting adjourned at 1130 on Wednesday 10th March.

Meeting convened at 1410 on Wednesday 10th March.

Revisiting some deferred comments…

Comment 213: Norman Finn

43.2.10 page 76 line 41 – change 'Identifier' to 'MAC address'. Replace all text and note to 43.2.11.2 (page 78 line17) with a description of the Aggregator Identifier, tying it to the ifIndex attribute. Add a sub-clause to explain Port Identifier.

Comment 270: Stephen Haddock

Change Aggregator_Port_Number to Aggregator Identifier. Simplify definition.

Change Aggregator to Aggregator_MAC_Address, page 91 line 1. Change identifier to MAC address.

Comment 271: Hadriel Kaplan

Fixed by 270.

Comment 256: Graham Short

No change to this line. See 43.2.10.

Comment 276: Graham Short

Change Aggregator to Aggregator Identifier.

Comment 332: Norman Finn

See 213.

Comment 50: Keith Klamm

Change to references to half-duplex. Link Aggregation will treat any half-duplex links as individual links and will not operate the Link Aggregation Control Protocol on these links.

Comment 365: Hadriel Kaplan

Modify page 109, line 54. State that half-duplex links shall be treated as individual with transmission/reception of LACPDUs "turned off".

1.4 State Machines

Peter Saunderson presented the results of the offline discussions on State machines (see Web site).

Receive Machine

Original states: CURRENT and EXPIRED supplemented by DEFAULTED, INIT and DISABLED. EXPIRED goes to DEFAULTED after a timeout.

Discussion about new Port_disabled state (disabled means link down or disabled by management). This is entered when a disabled event occurs, subsequently entering the INIT state, if the ports parameters are seen on another port. Transition from DISABLED to EXPIRED state after a port enabled event.

Must also enter DISABLED state if the port is half-duplex.

Periodic Machine

Only has NO PERIODIC, FAST PERIODIC and SLOW PERIODIC states. Start in NO PERIODIC, stay there until someone becomes Active.

Global entry into NO PERIODIC if half-duplex. Leave this as an action item to be determined off-line.

Combine all Global arcs entering NO PERIODIC state as a single arc.

Selection Machine

Virtually identical to existing document. Basically just wait a while before going to ATTACHING state, to allow other ports to ATTACH at the same time. The 'Ready' variable is replaced by the existing 'Selected' variable, with a value of 'SelectedSTANDBY'. There are numerous textual changes that must be made to describe this behaviour.

The above state machines, along with associated text changes, will be put into the document by the editors as a resolution to Mick Seaman's comments on the state machines.

1.5 Other Discussions

Some discussion took place on the need for key changes and the need for system and port priorities. The feeling is to leave this as is for now. If there is a need for this to change later, it is not envisioned to affect the state machines again.

This was based on Norman Finn's comments. Other related comments on items missing from management will be accepted. A better description, in the informative Annex, will be made to explain the rationale of the priorities and how they work. Comment 362. Priority mechanisms stay in the document. Change the name of the Annex, appropriately.

43.5.1 c) Clarify the highest priority port as being from the system of the highest priority.

Global search for uses of Key, to clarify whether it is Administration or Operational Key.

1.6 More Comments (Hot Editorials)

Comment 200: Hadriel Kaplan

Accept.

Comment 63: Vince Bridgers

Accept. Check these are defined according to the Table 30-2.

Comment 170: Hadriel Kaplan

Accept.

Comment 210: Keith Balmer

Accept. Change to 'a short time'.

Comment 245: Hadriel Kaplan

Reject.

1.7 Procedural Discussion

Stephen Haddock explained that this ballot was a Task Force ballot, hence informal instructions to the editors to deal with editorial issues, and even some technical issues, are okay. A Working Group ballot does not operate in this way, every comment must be strictly responded to with precise changes to the text.

Rich Seifert asked the groups approval to make the extensive changes that have come out from the comments and state machine changes, without the normal striked text indicating the changes from the previous draft. Also no change bars will be applied.

No-one objected to the editor producing draft 1.1 with the comment resolutions discussed above.

Interim Meeting

Favourite date is first week in May, co-located with 10 Gig group, in Boulder Colorado. This gives time for a new draft and Task Force ballot beforehand.

Minutes From Previous Meeting

Approved by acclamation.

Meeting adjourned at 15:20 on Wednesday 10th March.