UNCONFIRMED MINUTES
IEEE 802.3ad (Link Aggregation)
Task Force Plenary Meeting
November 9-13, 1998
Albuquerque, New Mexico

Prepared by:
Thomas Mathey

The meeting convened at about 9:10 AM, 9 November, 1998. Steve Haddock, Task Force Chair, opened the meeting with a presentation of the agenda. The following items were presented:

E-Mail reflector and website
the E-Mail exploder is stds-802-3-trunking@ieee.org
to subscribe, send E-Mail to: majordomo@majordomo.ieee.org
containing the line
subscribe stds-802-3-trunking@ieee.org <your E-Mail address>
the website is at: http://grouper.ieee.org/groups/802/3/ad/index.html

802.3ad Timeline with an emphasis on key dates
November 1998; cut-off date for proposals
January 1999; First Draft
March 2000; Standard

Objectives and presentations
"Big Ticket" Issues

Mr. Haddock then volunteered Thomas Mathey to act as recording secretary for the meeting. After introductions were made, Steve:
requested a show of hands for those people who were:
first time attendees
on the E-Mail reflector
aware of the previous papers which had been presented.

Steve then provided a brief recap for each of the "Big Ticket" Issues, followed by a detailed discussion of each issue.

Discussion: Specification of Distributor frame ordering constraints

Main points from discussion:

1. Steve provided a brief review of the reason behind the requirement; some protocols break if out-of-order frames are received. This discussion clarifies that all members are in agreement with the requirements listed in the LA Control Protocol.

2. No one distribution algorithm works for all applications/configurations, hence the difficulty in standardising this. However, frame ordering must appear to be maintained as far as the MAC service user is concerned.

3. An informative annex is extremely useful, and for the first draft, we need an outline.

Discussion: Architectural Model: # of aggregate ports required

Main points from discussion:

1. Two models (as shown in the agenda handout) are possible. For these models:
   A. are we in the low end, cheap markets and products, or
   B. are we at the higher end where customers want RAS (Reliability, availability, and Serviceable).

2. Should the standard require that all products have the number of PHY ports equal to the number of Aggregators. This is a “how tightly should the selection rules be for the binding of aggregator with PHY”.

Consensus:
1. The model is sufficiently general to describe all of the cases we have presently discussed. This side-steps the issue of “do we allow .....”.

2. The case of “the number of Aggregators is less then or equal to the number of PHY ports” is allowed.

Discussion: Addressing: MAC address per aggregate or per physical

Main points from discussion:

There are two possible views:
1. one from the top looking towards the bottom, which results in one (1) MAC address, and
2. one for the bottom looking towards the top, which results in one (1), (n), or (n+1) MAC address.

There may be a need to distinguish between/among the links for power on reset, debug purposes, and management requirements. Management, which may not be required by the standard, would then either require “n” addresses or a unique MAC address along with perhaps a port number.

Consensus:
1. Aggregator has one (1) MAC Address
2. PHY ports are only required to recognize a single MAC Address at a point in time
3. When bound, respond to aggregator
4. When not bound, expect a unique address for emitted packets
   A. It is not clear that this is an absolute requirement, but is better than any other solution; ie., what is the source address on LA Control packets.
5. The OUI is not required to be preassigned in a ROM, but can be assigned at a later time as long as it is unique.

Discussion: Dynamic Key modification rules

The main issue is: What are the rules from start at power-on-reset, receive a response (from the LA Control protocol ??), and then go on.

Main points from discussion:
1. We need some way to prevent an infinite loop
2. Parameters are available to management, and can only be changed by management
3. More flexibility than item 2 above may be needed to be useful and force/allow protocol to converge

Discussion: Selection rules

Main points from discussion:

wait for the presentation

Discussion: MAC addresses in LA Control frames

Main points from discussion:

There were good reasons from the Austin interim meeting for both a new address (from the reserved list) and a new Ethernet type field. The discussion then drifted in the next issue, Extensibility of LA Control frames.

Using a new address allows the address to have uses/extensions beyond Link Aggregation. Also, extend the type field by adding (op codes ??) for specific protocol uses. A 10 to 20 frames per second transmit with a 0.25 to 0.5 second response is suitable.

Since extension capability requires some up front work, Steve took a straw poll vote on how far shall we go in this “service to humanity”. General consensus was to go forward with “extension capability” and use due diligence. Now the issue is what to specify in the standard in order to prevent future “tales of woe”.

Discussion: Extensibility of LA Control frames

Main points from discussion:

The "extension" refers to the protocol itself, not to a future new protocol. (Also see previous discussion.)

Discussion: Managed objects for LA Control and "logical mac"

Main points from discussion:

MIBs

1. The 802.3 standard uses GDMO, but the rest of the world (ie., the IETF) uses SNMP. To drop GDMO would require that the Link Aggregation task force go to the 802.3 working group for advice and concent on adding in SNMP and/or dropping GDMO.

Note from the recording secretary: When this subject was brought before the working group, my memory says that dropping GDMO and using SNMP was OK, even if this meant that the standard is now written in two styles. The advantage of not having two MIBs and not requiring a translation from GDMO to SNMP out weighed the disadvantages.

2. The logical MAC needs a suitable set of full-duplex objects.

3. A list of MIBs:
   A. Interface MIB
   B. Ethernet like objects
   C. LA Control Protocol

4. Anil R. was one of several volunteers to provide and work on the MIBs.. I think that David Law volunteered some web site activity.

Manual Configuration

The issue of manual configuration and its relationship to the protocol and MIB variables was then discussed.

1. Is the LA Control protocol required.
2. Is manual configuration to be allowed.
3. Is forced settings of MIB variables to be allowed.
4. Can the LA Control protocol overwrite the desired settings of the MIB variables.
5. Not speak the LA Control protocol unless spoken to.
6.. Can manual configuration use the LA Control protocol.

Discussion on the Manual Configuration issues was postponed until after the two presentations.

End of task force meeting for Monday, 9 November 1998; 9 AM to 12 PM.

Start of task force meeting for Tuesday, 10 November 1998; 1 PM to 5:15 PM

The meeting convened at about 1:00 PM, 10 November, 1998. Steve Haddock, Task Force Chair, opened the meeting with a recap of the agenda, presentation handouts, timeline, attendance book, and the search for an interim host in January.

The meeting minutes for both the La Jolla plenary and the Austin interim were accepted by acclamation.

Steve then presented the call for patents talk, per the Geoff Thompson style as follows (following text copied/cribbed directly from the very excellent Merrimack minutes):

Steve, as the virtual Mr. Geoff Thompson, described the IEEE patent policy, and requested that all holders of patents related to the work of the study group become familiar with the policy and submit letters if appropriate. He also informed the group of the 802.3 policy to also request letters for patent applications. A description of the policy and example letters can be found at the IEEE web site (http://grouper.ieee.org/groups/802/3/patent.html).

Presentation: LA Control Protocol Update

by Tony Jeffree
from 1:15 to 4:10

Presentation: Selection Rules

by Jeff Lynch
from 4:15 to 5:15

Main points from discussion:

1. The main delta between Tony Jeffree presentation and the Jeff Lynch presentation is that this proposal allows for increased determinism vs the Tony proposal for “Plug & Play”. Of course, this lead immediately into : “please define determinism” along with which proposal will show/detect an error.

2. For selection rules, the “core” of a network wants extreme control. For example, after a power outage, the box or boxes will power on reset to a known, stable, identifiable configuration with “all” parameters/configuration variable, IP addresses in a filter table and assignment to AgPorts, all AgPort vs PHY port assignments remain the same, etc.

3. For selection rules, the “edge” of a network wants flexibility and “Plug & Play”. For example, a box establishes connectivity even if the resulting configuration is different from previous configuration (and ?? different from requested configuration via MIB variables).

4. There are cases where the higher layers want to know about lower layer connectivity changes. If there were no selection rules in the standard, then we could end up with an incompatible result between various vendors products.

5. there is value in assigning a key to an AgPort

Take a vote on Wednesday

End of task force meeting for Tuesday, 10 November 1998; 1 PM to 5:15 PM

Start of task force meeting for Wednesday, 11 November 1998; 10:25 AM to 3:00 PM

The meeting convened at about 10:25 AM and Steve Haddock went straight into discussions.

Discussion: Flush rules

Main points from discussion:

1. The better the flush (command ??) works, then the less need there is for a flush response (??). Any use of FLUSH is an improvement.

2. The flush command should have a simple return/response along with a default time-out of one second which can be changed by the user via a management parameter.

3. Steve then took a straw poll for the following motion:
   A. Response time limit for the collector responding to a FLUSH request
      YES: 3   NO: 30   Abstain: ?

   B. For distributor, specify a time-out for distributor for waiting for a flush response
      YES: 13   NO: 25   Abstain: 35

   Tony Jeffree then asked for the following change to the LA Control protocol which was approved without contention:

"remove response field from return"

Discussion: Determinism

Main points from discussion:

Steve proposed the following levels, which were modified during the discussion

1. After a power-on-reset for a given static state for cabling, arrive at a repeatable state
2. After a (arbitrary) sequence of changes, [repeat the same sequence], arrive at a repeatable state
3. ~~After reset, arrive at same state as before reset (history independent)~~
   State is only dependent upon physical cable configuration)

A paradox is possible depending upon choice of state vs sequence. The standard may need to impose a time limit after power-on.

Discussion: Selection Rules

Main points from discussion:

There are several options possible.

1. Is there any need to specify the selection rules. If specify, is there one set of rules or multiple sets. There may need to be different sets of rules for different types of products.. It is possible to offer one (1) rule as an example with a requirement that all products meet a specified behavior.. One person voted to keep "Plug & Play".

2. There is no one (1) set of rules which meets all needs. Thus there are no normative rules, and this is a good subject for an informative annex.

3. One issue worth doing is "Plug & Play", and then use network management for those users who know what they are doing.

Steve summarized discussion as: allow flexibility but provide a default; much like present day bridges do. Possibly recommend a default set. Jeff Lynch provided a 1 page summary of rules repeated below:

Selection Logic

- General Rule
  - Each physical port is assigned a Key
  - Each AgPort is assigned a Key
  - Matching the physical port and AgPort keys determines trunk groups
    + If a link is not to aggregated, its port is given a unique key that matches its aggregator

- In the case of ambiguous AgPorts (i.e., more than one AgPort has the same key) include a Recommended (non-normative) default algorithm in specification
  - other example algorithms may also be included in an informative annex

---

As a result of discussion, the second bullet was dropped, but the text was kept for inclusion in the first draft.

Presentation: Dynamic Port Keys

by Norm Flynn
from 1:20 to 2:00

Main points from discussion:

1. Separate the Administration State/Key from the Operational State/Key. The Administration State is set to control or provide a "this is what I want" vs the Operational State of "this is what I got".

2. For a fixed configuration, then turn off LA Control Protocol, not override the protocol.

Steve then proposed the following motion:

Incorporate the Norm Flynn proposal/presentation for “Dynamic Port Keys” in the first draft.

YES: 40 NO: 0 Abstain: 2

Discussion: Management

Main points from discussion (led by David Law):

1. A proposal is to start with SNM, then convert to GDMO, and include both in the standard.
2. Editor to create without a presentation.
3. A 50% vs 75% vote for changes.
4. Send submittals to David Law for placement on web site; submittals are expected to be large and will not be sent out on the E-Mail reflector.

Discussion: Extensibility of LA Control frames

Main points from discussion:

1. Are buffer sized to be fixed size, and if so, then at what size?
2. Are unused fields to be padded/set to 0’s?

Discussion: Editorial Policy

Rich Seifert provided a review of how to obtain from the IEEE the MAC Address and Type fields that the LA Control protocol will require. These fields are not made public until sponsor ballot.

Tony Jeffree will modify his proposal to meet 802.3 formats for state machines, for MIB, and for service interfaces. There will be changes to existing clauses, along with a new clause 43. Draft 0.1 is expected prior to the January interim, with editorial direction for a draft 1.0 coming out of the January interim.

Motion: First Draft

The task force charges the editor to generate draft version 0.1 of 802.3ad for distribution prior to and submission at the January interim meeting, based on this week discussions, presentations, and resolutions at this meeting

Moved: Tony Jeffree Second: Rich Seifert

YES: 38 NO: 0 Abstain: 0

Interim Meeting: Location and Date

To allow co-location with other 802.3 groups, this will be discussed at the 802.3 level Thursday morning.

Note from the recording secretary: At the closing 802.3 Plenary on Thursday morning, Level One offered to host the interim meeting at a location in the south Florida area during the week of Jan. 18-22.

End of task force meeting for Wednesday, 11 November 1998; 10:25 AM to 3:00 PM

Note from the recording secretary: The large meeting room was fairly full, and while I neglected to count, well over 100 attendees were present.