UNCONFIRMED MINUTES

TRUNKING STUDY GROUP

Merrimack, NH
28-29 Apr 1998

TUESDAY, 28 APR 1998

INTRODUCTORY REMARKS AND INTRODUCTIONS

Mr. Steve Haddock called the meeting to order at 9:00. He invited participants to introduce themselves. Mr. Haddock also thanked Mr. Ben Brown of Cabletron for hosting the meeting and for the fine accommodations and toys (a nice projector). Mr. Haddock then volunteered Mr. Robert Grow to act as meeting secretary.

IRVINE MEETING REVIEW

Mr. Haddock briefly reviewed the Irvine meeting and indicated that the minutes of the meeting would be considered for approval on Wednesday.

PRESENTATIONS

Mr. Paul Bottorff of Bay Networks make a presentation on "Link Aggregation Configuration". The presentation discussed some issues between single and multiple client architectures. Mr. Bottorff proposed the standard should use a single client architecture. He also presented some interaction issues between aggregation capable and aggregation non-capable DTEs.

The binding of MACs to a MUX was also discussed along with the general structure of MAC Aggregation Control. In discussion it was agreed that portions of MAC Aggregation Control had to exist with each MAC to allow any MAC to be the first in an aggregation group. The extent to which another part of MAC aggregation is shared between all MACs in the group is dependent on the features placed within the sublayer.

Mr. Bottorff concluded by proposing a basic three state machine to handle non-capable devices, and aggregation capable devices, with the characteristics of the states.

Mr. Tony Jeffree presented on "Link Aggregation Control Protocol", describing a document he distributed on the reflector prior to the meeting. He included his basic assumptions and objectives for aggregation. Multiple comments were made on his determinism objective and what was included in it. Mr. Jeffree defined determinism as the behavior that independent of the order in which links activate the aggregation group will always consist of the same links. This did not necessarily include the addressing of the aggregation group.

Identification of link characteristics and groups supports the automatic nature of the group control. Mr. Jeffree proposed that identification be a local system ID and capability ID. The combination of identifiers at both ends provides the necessary uniqueness for defining an aggregation group.

To prevent duplication and reordering problems, a set of rules was proposed. He also proposed basic protocol rules and operational considerations. The valid link states were a point of considerable discussion. This included the sequencing of starting collection and distribution, semantics of flushing and other fault requirements.

Mr. John Wakerly, Mr. Michael Fine of Cisco Systems presented on "Link Aggregation Protocol Techniques", also authored by Mr. Norman Finn. The presentation summarizes things learned in existing proprietary link aggregation products. Mr. Wakerly gave a summary of link aggregation requirements and the overall scheme for managing aggregation.

Mr. Fine described details of the per-link state machine. It supports link aggregation capable and non-capable devices. The discussion included attempts to correlate this state machine with Mr. Jeffrey's link states, though there was no direct correlation. The per-link state machine is designed to achieve stability before initiating MAC client traffic. This minimizes thrashing as links come up.

The link group is controlled by an aggregation monitor. The available links are determined by identifiers and the per-link state. The identifiers are similar to those proposed by Mr. Jeffree consisting of a device identifier and a locally assigned category number (similar to SystemID and CapabilityID), but used for equality comparison by the aggregation partner.

A document, "Port Aggregation Protocol", detailing the proposal was distributed and will be available on the web.

DISCUSSION ON LINK AGGREGATION CONTROL

Mr. Haddock noted that this had been the topic of the presentations. There was little more to discuss, beyond the presentation discussions.

WORK PLAN

June: Task Force becomes official

July: (Plenary) and September (Interim):
* Call for proposals
* Refine proposals on the table

November (Plenary):

* Last opportunity for proposals
* Select proposal to create first draft
* Possibility of coming into Nov. meeting with first draft.

CALL FOR PATENTS

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).
Plan for next meeting

GENERAL DISCUSSION

Mr. Haddock led the discussion on open issues. The following list was composed.

* Management
  - New counters?
  - Status indications
  - Control functions/parameters
* Frame format
  - MAC control versus a new Ethertype
  - Addressing
  - Contents and format
* Relationship between MAC Control and Link Aggregation Control
  - Document structure
  - Interaction with flow control
* How do we deal with dynamically creating/removing MAC Client interfaces
  - Effects on spanning tree protocol, routing protocols
  .-.How does link status get reported to MAC Clients
* Does link aggregation have to stabilize before enabling data traffic? Is some traffic handled different than others, e.g., spanning tree?
* Do distribution rules need special handling for certain types of traffic e.g., spanning tree protocol?
* Flush requirements
* Handling packets during LAG MUX transitions
* Is there a need for an upper limit on the number of links that can be aggregated?
* Deterministic configuration
  - Time invariant topology resolution
  - Determinism of address binding
  - Determinism vs. stability
* Protocol needs a periodic keep alive for aggregations.
  - At what interval?
* Interaction of manual configuration with auto configuration.

- Can a protocol aware station create a group with a manually configured non-protocol aware station?
- Can manual configuration override auto configuration but still use protocol for verification of configuration?

* Protocol should handle removal of links from aggregate whether intentional or failure at link.
  - How quickly
  - At/on suspicion of connection verses assurance of connection?
* How do we define link aggregation sublayer interfaces?
  - Incorporate 802.1D internal sublayer service?

These issues were then grouped:

1) Transparency issues
   - Dynamic addition/removal of MAC clients
   - Impacts on higher layer protocols
2) Protocol issues
   - Determinism vs. stability (or selected by management policy)
   - Flush
   - Keep alive
   - Link removal, failure
   - Interaction with manual configuration
3) Structure of standard
   - Interface definition
   - Relationship to MAC control
4) Management

A lengthy discussion followed on the architectural model for handling additions and deletions of MACs to aggregation groups. Alternate architectures were described. While there was some disagreement on functionality, the bulk of the discussion was to clarify what was included within a architectural component, and how they are instantiated. The impact on the structure and content of the standard was also a frequent consideration. The discussion drifted over many topics including address binding, MIB instantiations, use of existing MIBs, and the appearance of an aggregation group to existing management applications.

The discussions produced the following general agreements of things that were expected to be approprate in an aggregation standard.

* The draft standard should deal with one link aggregation (LA) control entity. Each LA has an ID (DeviceID in presentations).

* The LA Control entity may contain multiple capability groups (potential LA groups).

* An LA Control entity has one or more interfaces to physical MACs.

* An LA Control has one or more MAC Client interfaces (up to the number of physical MACs?)

* A MAC Client interface has a permanently bound MAC address.

Mr. Haddock took the action item to summarize the discussion and agreements for the group Wednesday morning.

WEDNESDAY, 29 APR 1998

Mr. Haddock called the meeting to order at 9:00 am. He presented his summary of the discussion and points of agreement from yesterday's discussion. The presentation as clarified in by the group will be posted to the web.

The discussion was focused on a more detailed diagram of the link aggregation sublayer. There was general consensus that the aggregator (a MUX in most presentations) contains the collector and distributor function where multiple MACs are presented to the MAC Client as a single logical MAC. There was also general consensus that there is some state and perhaps a state machine associated with each physical MAC. The bulk of discussion was on the control function.

One view is that there is a single link aggregation control spanning all aggregators in a device. The alternate view is that the control function is not necessary in that the aggregator and per MAC state machine could sufficiently specify the protocol. The consensus of those commenting was to minimize the control function and that the more detailed diagram presented by Mr. Haddock is intended to stimulate more consistent use of terminology in presentations as the focus for discussion, not for inclusion in the standard.

The group then revisited the addressing issues, including the assumption that the MAC address is bound to the aggregator. Some participants contended that initially each aggregator would get an address from a default MAC. It was pointed out that the Irvine meeting had concluded that each link aggregation sublayer had a unique address. This correlates to aggregator within the more detailed architecture, which during this meeting was also assumed to have a unique address.

MINUTES FROM IRVINE MEETING

The minutes of the Irvine meeting were corrected and approved by voice vote.

Mr. Haddock reviewed the work plan indicating where work was needed, and presentations topics that would be welcomed at the July meeting in San Diego.

UPCOMING MEETING

The study group is anticipated to meet as a task force at the July meeting, and requested of the 802.3 chair that it be scheduled for meetings on Tuesday and Wednesday. If an 802.1 plenary is held on Wednesday, the link aggregation meetings would be adjusted to allow plenary attendance.

With no further business to conduct and without objection, the meeting was adjourned at 12:00 noon.

Link Aggregation Control Assertions

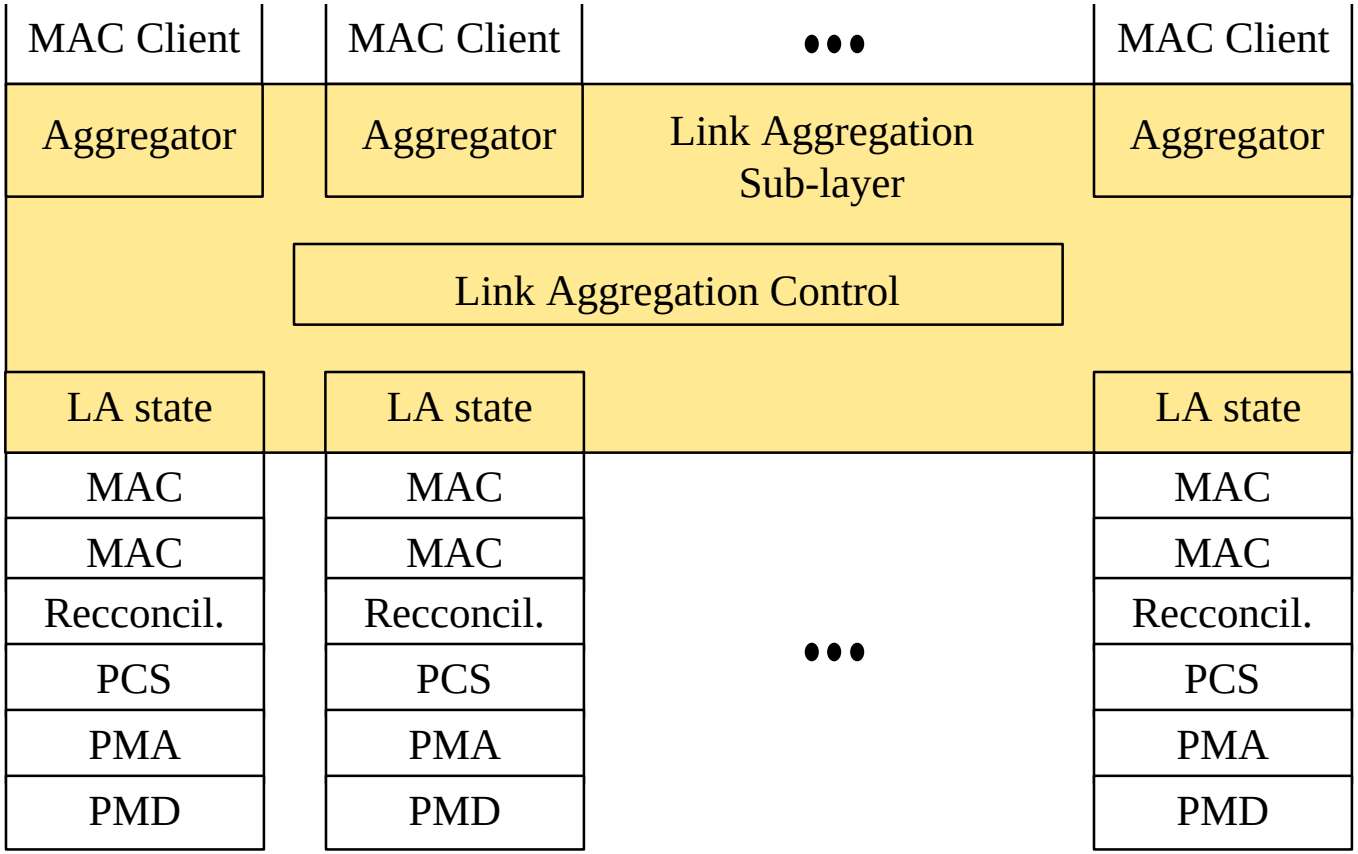

Assertions -- LA Control

- Link Aggregation sub-layer assumes one LA Control entity
  - An LA sub-layer has a single unique ID (called a System ID or Device ID in the protocol presentations).
  - LA sub-layer has one or more interfaces to physical MACs (interface is to the top of the MAC Control sub-layer).
  - LA sub-layer contains one or more Aggregators.
  - The function of the LA Control is to coordinate the linkage of zero or more MACs to each Aggregator.
  - LA Control associates a Capability Group (aka Category) number with each MAC. There may be multiple groups/categories within a single LA Control entity.

Assertions -- Aggregator

- There is one Aggregator for each MAC Client supported by the device.
  - An Aggregator presents a single MAC Client interface to higher layers.
  - Each Aggregator has a permanently bound MAC address.
  - Each Aggregator contains a single Link Aggregation Collector and Distributor function.
  - An Aggregator appears to be a single MAC to all higher layer functions connected at the MAC Client interface.
    - An Aggregator that is not linked to any MACs appears to the MAC Client to be a single MAC on an inactive link.

Assertions -- Addressing

- MAC Addresses are permanently bound to a MAC Client interface (i.e. to an Aggregator).
  - One-to-one mapping of MAC Client interfaces to Aggregators.
  - Requirements for the address of each MAC Client interface to be unique are no different than in current systems with multiple MAC Clients where there is a one-to-one mapping of MAC Client interfaces to physical MACs.

Assertions -- Addressing (cont.)

- There is no clear requirement for physical MACs to be assigned unique addresses.
  - For all frames other than MAC Control and Link Aggregation Control frames (ie all frames traversing a MAC Client interface), the physical MAC takes on the address of the Aggregator to which it is connected.
  - The address used in the SA field of MAC Control and Link Aggregation Control frames is TBD, but may not need to be unique.
  - Implementers are not precluded from assigning a unique address to each physical MAC.
  - In several presentations it has been requested that NICs in end stations not be required to operate with more than one individual address, in which case the one would be the MAC Client interface address.