

# Things that go Bump in the Night

Mick Seaman

This note describes one way to extend the proposed standard Link Aggregation Control Protocol to handle “full duplex shared media repeaters”. These have not been standardized by IEEE 802. However there is a concern that such devices may be encountered in real world networks.

If they do, and if there is to be no operational requirement for manual configuration of LACP, it would be as well to be able to accommodate such devices gracefully. It is a goal to allow LACP deployed everywhere by default, without user resistance arising from undesirable behavior in cases where it is not particularly wanted, so the protocol should at least be capable of using a link to a shared media full duplex repeater as a normal “individual” non-aggregated link.

There is no suggestion that the extensions described here should be standardized. However ensuring that such extensions are possible is “due diligence” against future standardization of the full duplex buffered repeater. A system implementing any such extensions needs, of course, to function in accordance with the proposed 802.3ad standard protocol where a point to point link is truly involved.

## Goals

This note puts forward a way or ways of extending the proposed 802.3ad link aggregation control protocol to deal with the case where a full duplex link attaches to shared media and the attaching DTE is unaware from physical signaling that shared media and not a point to point link is involved.

The goal is to support correct operation of a switched LAN where three or more bridges and endstations are attached to one or more instances of such shared media in the switched LAN and all these bridges and endstations are running LACP – though none of the mysterious full duplex repeaters are.

It is a non-goal to ensure correct operation where only some of the bridges attached to such shared media are running LACP. The author has no solution to offer for the case where two bridges are running LACP, and can therefore include the shared media link segment as part of an aggregate, while further bridges not operating link aggregation are attached to the same shared media.

The objective of the proposal is to ensure that the shared media link continues to be used as an individual link, i.e. to defend against the following undesirable outcomes when there are more than two protocol partners attached to the shared media:

- a) continual protocol thrashing due to repetitive change of partners means the link is unusable
- b) the link intermittently comes into operation as part of an aggregate involving only two of the partners.

One way of using shared media in an aggregated link would be to aggregate together links with the same “name” where the name is the set of system id/key id pairs of the devices

attached to the link segments. This would involve each attached system keeping track of a number of concurrent protocol partners. While this is certainly possible, that is not what is being proposed here<sup>1</sup>.

Rather the goal is to ensure that a shared media link will always be treated as an individual link if more than two LACP protocol partners (one of whom at least is active) are attached to the shared media.

## Extensions

The extensions are the addition of a Crowd Control Machine, similar to that described in initial protocol proposals, and modifications to the Periodic Transmission Machine.

The Crowd Control Machine keeps track of partner changes, and if two or more such changes occur in a Slow Timeout interval, classifies the link as an “individual link” for the Selection Logic, and sets the Aggregate(able) flag transmitted as part of the actor’s State to be false. That ensures all protocol partners will treat the link as an individual link.

Additionally the Periodic Transmission Machine maintains itself in the Fast Periodic state if any participant (not just the last PDU received) has communicated Short Timeout required in the last Slow Timeout interval and there is any active participant. Again the test for active participants refers to the actor and all LACPDUs received in a preceding Slow Timeout interval – not just the last LACPDU received.

These extensions are believed sufficient to meet the goals described.

---

<sup>1</sup> Even if the existence of shared media were to be admitted into 802.3ad such a solution would still be outside the scope of the standard.