+---+
| IEEE 802.1 REVISION REQUEST 0009 |
+-----------------------+=================================-----+

DATE: September 8, 2011
NAME: Jeffrey Lynch
COMPANY/AFFILIATION: IBM
E-MAIL: jjlynch@us.ibm.com

REQUESTED REVISION:

STANDARD: IEEE 802.1Q &/or IEEE 802.1AX
(possibly included into the 802.1AXbq (Distributed LAG) specification/project)

CLAUSE NUMBER: Appropriate clauses for LinkLAG MIB and TLV definitions such as P802.1Q-REV Annex D "D.2.7 Link Aggregation TLV"

CLAUSE TITLE: Appropriate clauses for LinkLAG MIB and TLV definitions such as P802.1Q-REV Annex D "D.2.7 Link Aggregation TLV"

RATIONALE FOR REVISION:

The current IEEE 802.1 standards are silent on how LLDP is to operate over an aggregated link. This could lead to incompatible implementations in the marketplace.

- Is there one LLDP session per physical link or one for the aggregated link? Or are there both?
- Background
  - new-pelissier-laglldp-0110.pdf "Notes from LLDP over LAG concall"
  - new-congdon-linkag-LLDP-0110.pdf "Disambiguating LLDP agents over a Link Aggregation"
  - new-lynch-LLDP-over-LAG-0911.pdf "Eliminating the ambiguity of running LLDP over LAGs"

PROPOSED REVISION TEXT: (See background presentations above)

1) Document in IEEE 802.1Q or IEEE 802.1AX the requirements and operational behavior for LLDP in the present of Link Aggregation
   - Either redefine the use of some of the reserve bits or add new fields to the existing LAG TLV (to disambiguate the source & target of LLDP frames)
2) Identify which TLVs should be sent on each individual link or on the aggregate
3) Resolve the architectural addressing issues introduced by running LLDP over a LAG.

IMPACT ON EXISTING NETWORKS:

None known at this time.
See background presentations (above) for potential backwards compatibility implications.

+---+

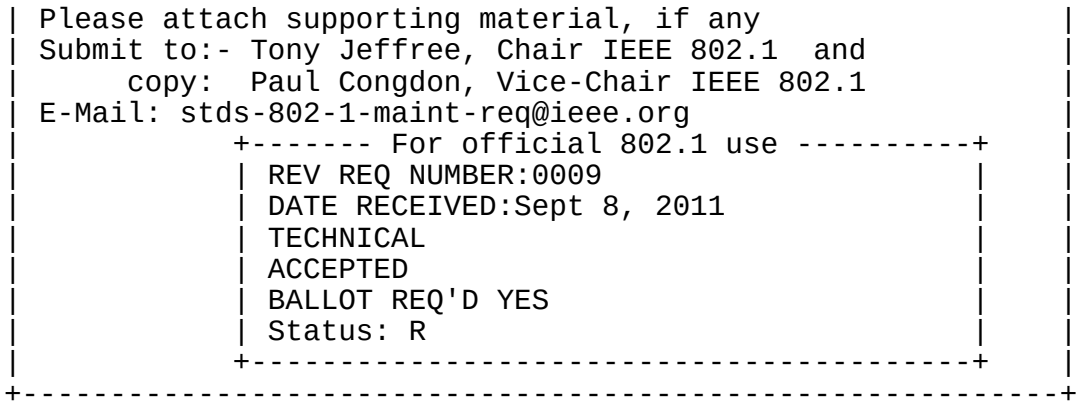
Please attach supporting material, if any
Submit to:- Tony Jeffree, Chair IEEE 802.1 and
copy: Paul Congdon, Vice-Chair IEEE 802.1
E-Mail: stds-802-1-maint-req@ieee.org
For official 802.1 use
REV REQ NUMBER:0009
DATE RECEIVED:Sept 8, 2011
TECHNICAL
ACCEPTED
BALLOT REQ'D YES
Status: R

Notes from LLDP over LAG concall

Joe Pelissier

With many thanks to Anoop Ghanwani for sharing his excellent notes ☺

January, 2010

new-pelissier-laglldp-0110

LLDP & Link Aggregation

- Background

  In the November IEEE 802.1 meeting, it was clear that there was confusion over how LLDP is to operate over an aggregated link

    Is there one LLDP session per physical link?

    Is there one LLDP session for the aggregated link?

    Are there both?

- Joe Pelissier agreed to arrange an informal conference call of interested parties to discuss

  This is a summary of that call

LLDP & Link Aggregation

- Conesus:
  - LLDP must be able to run on individual links.
  - LLDP may run over the aggregate.
  - Necessary to determine topology with link aggregation through TPRs
    - Norm to provide some material on this
  - We need a way to disambiguate the aggregate
    - Paul will have a contribution on this
- For DCBX:
  - For each TLV, we need to figure out whether to run it on individual links or aggregate.
    - Also need to determine convergence in case of a conflict
- May need to look at other TLVs, e.g. VLAN TLV
  - Does it makes sense to send it on the individual link or the aggregate.
- Manoj brought up the issue of end station problems with link aggregation
  - Not LLDP specific just general problems such as stateful offloads and link aggregation
    - Separate problem from the one under discussion here.

Disambiguating LLDP agents over a Link Aggregation

Paul Congdon

Objective

- Provide a mechanism for receiving LLDP agents to determine the source of LLDP frames over a link aggregation
- Propose a way to document this within our standards under construction

Interpretation of current specs

- 802.1AB does not say anything specific about how agents work in an aggregation environment
- The current LinkAgg-TLV has fields to indicate capability of being aggregated and current status of the aggregation
- Assumption from 802.1AB is that agents would exist on the physical links and potentially on the aggregation as well

Example LLDP over LinkAgg

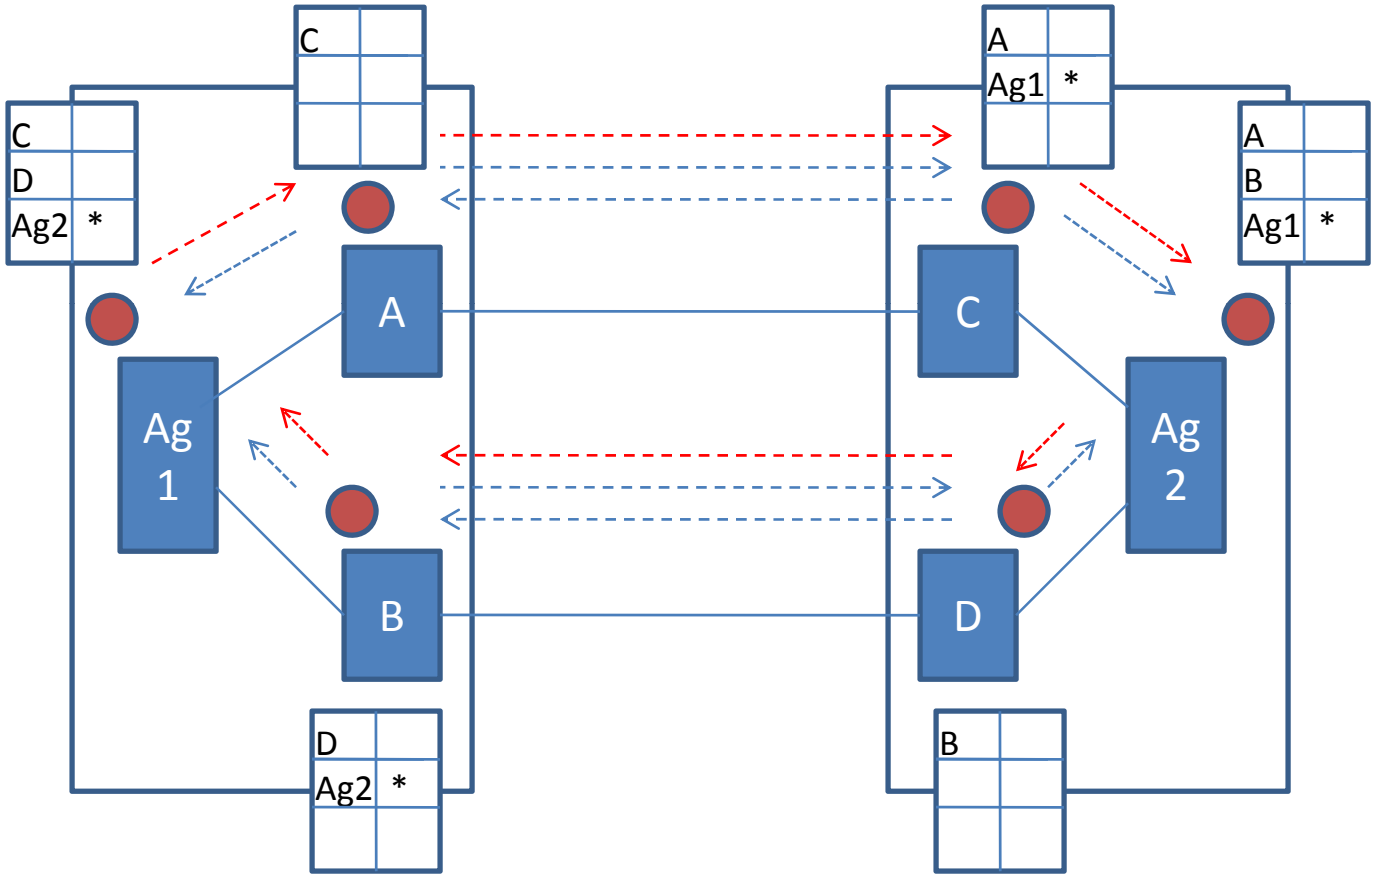

Note: Aggregation agent LLDP frames are distributed over the aggregation to one link
Aggregation agent sees all LLDP agents
Physical link agent sees physical link peer and possibly the aggregation agent

Disambiguating Agents

- Require LinkAgg TLV to be present if running in a Link Aggregation

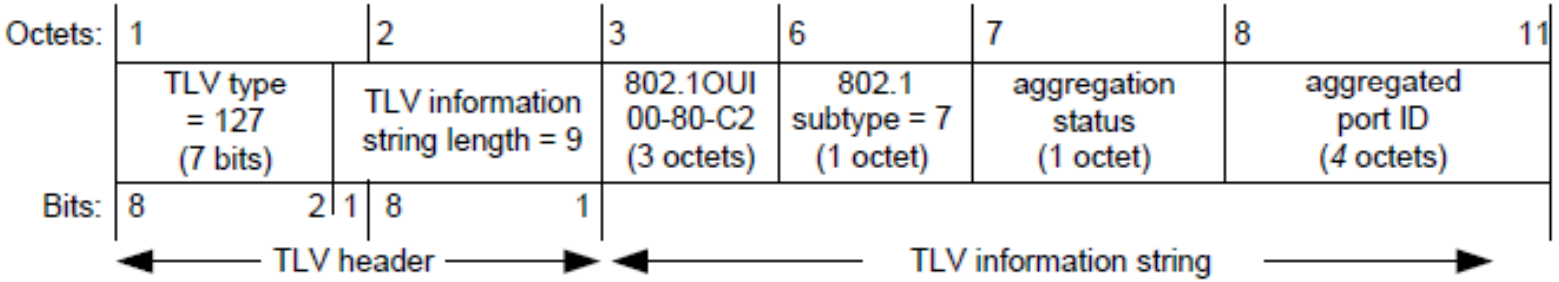

Figure E.7—Link Aggregation TLV format

Table E.3—Link aggregation capability/status

Bit	Function	Value/meaning
0	Aggregation capability	0 = not capable of being aggregated 1 = capable of being aggregated
1	Aggregation status	0 = not currently in aggregation 1 = currently in aggregation
2–7	reserved for future standardization	—

Disambiguating Agents (cont)

- Agents on physical links
  - Aggregation capability = 1
  - Aggregation status = 1
  - Aggregated PortID = PortID of aggregated link
  - PortID TLV = physical link PortID
- Agents on aggregates
  - Aggregation capability = 0
  - Aggregation status = 1
  - Aggregated PortID = PortID of aggregated link
  - PortID TLV = aggregated link PortID

Standardizing this proposal

- There is a proposed PAR to open 802.1AX to document how to use LinkAgg over TPMRs
- Insert a small clause to document requirements for LLDP in the present of LinkAgg

Issues

- Will setting capable=0 and status=1 trigger error conditions within existing management stations?
- Aggregation portID can not be the same as any physical portID for this to work.
- Chassis-ID+Port-ID must be unique for each agent to be identified. If the aggregate uses the same Chassis-ID+Port-ID the receiver will wipe-out previously received data.
- This proposal makes it impossible to aggregate an aggregation. This may be already true
- Adds LLDP to the long list of protocols that must be coordinated in the implementations of Link Aggregations split across multiple chassis.

Alternatives Discussed

- Define a new address to direct frames only to the aggregate (e.g. Nearest-Aggregate Address)
- Create a new mandatory TLV that must only be sent by the aggregate agent

Eliminating the ambiguity of running LLDP over LAGs

(Revised based on feedback during 9/6/2011 EVB conference call)

Jeff Lynch

September 6, 2011

Background

- In the November 2009 IEEE 802.1 meeting, the issue was raised concerning how LLDP is to operate over an aggregated link
  - Is there one LLDP session per physical link?
  - Is there one LLDP session for the aggregated link?
  - Are there both?
- Particular clarification is needed in the context of DCBX

LLDP & Link Aggregation

Consensus of January, 2010 telcon (documented in new-pelissier-laglldp-0110)

- LLDP must be able to run on individual links.
- LLDP may run over the aggregate.
- Necessary to determine topology with link aggregation through TPMRs
  - Norm to provide some material on this (Contribution not yet submitted)
- We need a way to disambiguate the aggregate
  - Paul will have a contribution on this
  - Contribution submitted: new-congdon-linkag-LLDP-0110.pdf "Disambiguating LLDP agents over a Link Aggregation"
- For DCBX:
  - For each TLV, we need to figure out whether to run it on individual links or aggregate.
  - Also need to determine convergence in case of a conflict
- May need to look at other TLVs, e.g. VLAN TLV
  - Does it makes sense to send it on the individual link or the aggregate.
- Manoj brought up the issue of end station problems with link aggregation
  - Not LLDP specific just general problems such as stateful offloads and link aggregation
  - Separate problem from the one under discussion here.

Text in blue added to represent current status

LLDP & Link Aggregation

Review Paul Congdon's slides: "Disambiguating LLDP agents over a Link Aggregation"

new-congdon-linkag-LLDP-0110.pdf

IEEE 802.1 Revision Request

```
+-------------------------------------------------------------------+
| IEEE 802.1 REVISION REQUEST XXXX                                  |
+------------------=================================---------------+
DATE:    September 6, 2011
NAME:    Jeffrey Lynch
COMPANY/AFFILIATION:  IBM
E-MAIL:  jjlynch@us.ibm.com

REQUESTED REVISION:
     STANDARD:  IEEE 802.1Q &/or IEEE 802.1AX (possibly included into the 802.1AXbg (Distributed LAG) specification/project)
     CLAUSE NUMBER: Appropriate clauses for LinkLAG MIB and TLV definition
     CLAUSE TITLE:  Appropriate clauses for LinkLAG MIB and TLV definitions

RATIONALE FOR REVISION:
    The current IEEE 802.1 standards are silent on how LLDP is to operate over an aggregated link.  This could lead to
           incompatible implementations in the marketplace.
        •     Is there one LLDP session per physical link or one for the aggregated link?  Or are there both?
        •     Background
              •     new-pelissier-laglldp-0110.pdf “Notes from LLDP over LAG concall”
              •     new-congdon-linkag-LLDP-0110.pdf “Disambiguating LLDP agents over a Link Aggregation”
              •     new-lynch-LLDP-over-LAG-0911.pdf “Eliminating the ambiguity of running LLDP over LAGs”

PROPOSED REVISION TEXT:  (See background presentations above)
1) Document in IEEE 802.1Q or IEEE 802.1AX the requirements and operational behavior for LLDP in the present of Link
       Aggregation
       Either redefine the use of some of the reserve bits or add new fields to the existing LAG TLV to disambiguate the
             source & target of LLDP frames
2) Identify which TLVs should be sent on each individual link or on the aggregate
3) Resolve the architectural addressing issues introduced by running LLDP over a LAG.

IMPACT ON EXISTING NETWORKS:

    None known at this time.  See background presentations (above) for potential backwards compatibility implications.


+-------------------------------------------------------------------------------------------------+
| Please attach supporting material, if any                                                       |
| Submit to:- Tony Jeffree, Chair IEEE 802.1  and copy:  Paul Congdon, Vice-Chair IEEE 802.1      |
| E-Mail: stds-802-1-maint-req@ieee.org                                                           |
|           +------- For official 802.1 use ----------+                                         |
|           | REV REQ NUMBER:                          |                                         |
|           | DATE RECEIVED:                           |                                         |
|           | EDITORIAL/TECHNICAL                      |                                         |
|           | ACCEPTED/DENIED                          |                                         |
|           | BALLOT REQ'D YES/NO                      |                                         |
|           | Status: X                                |                                         |
|           +------------------------------------------+                                        |
+-------------------------------------------------------------------------------------------------+
```