

802.1AX -- Link Aggregation:

Editor's Report: January 2018

Version 1

Stephen Haddock
January 24, 2018

Agenda

1. Draft status and changes implemented as a result of discussions during 802.1AX Editor's Report in Orlando.
2. Proposal for support of three Portal Systems in a Portal.
3. Mick's proposed refinements of LACP state machines.


<http://www.ieee802.org/1/files/public/docs2018/ax-seaman-lacp-suggestions-0118-02.pdf>

Preparing 802.1AX-Rev-d0.2

- AX-Rev-d0.1 went to Task Group ballot in June and comment resolution in July
 - Major changes were in Clause 6 (LACP) and specifically relating to Conversation Sensitive Collection and Distribution (CSCD).
 - Clause 9 (DRNI and DRCP) was mostly untouched.
- AX-Rev-d0.2:
 - Incorporates comment resolutions from first Task Group ballot.
 - Major changes in Clause 9 (DRNI and DRCP)
 - The rest of this presentation summarizes those changes
 - No changes yet to MIB or PICS
 - Do this for Working Group ballot
 - Plan to have second Task Group ballot shortly after ~~November~~ **January** meeting (?)

Review of November's Editor's Report

1. Detecting Portal Topology Errors
 - Particularly with three Portal Systems.
2. Communicating Portal State
 - Particularly with three Portal Systems.



Comments from 802.1AX-Rev Editor's Report in Orlando

1. Should make the sequence number cover any state variables sent that need to be acknowledged, not just the Gateway_Enable_Vector.
 - Have separate sequence numbers for Gateway and Aggregator info.
 - Create Gateway_State and Aggregator_State structures that include the sequence number and other variables.
 - Use these structures in state variables and DRCPDUs TLVs
2. A Portal System at the end of a chain needs a sequence number acknowledgement from the Portal System at the other end of the chain, not just from the Neighbor.
 - Create a Portal_State_Summary structure that includes just the Gateway and Aggregator sequence numbers for each Portal System.
 - Use these structures in for acknowledgements in DRCPDUs.

Structures

DRF_Gateway_State
Gateway_Sequence_Number
Gateway_Enable_Vector
Drni_Gateway_Algorithm
Drni_Gateway_ServiceID_Map_Digest
Drni_Gateway_PortList_Digest

DRF_Aggregator_State
Aggregator_Sequence_Number
Active_LAG_Links
Actor_Port_Algorithm
Actor_Conversation_Service_Mapping_Digest
Actor_Conversation_LinkList_Digest
Partner_System_Priority
Partner_System
Partner_Oper_Aggregator_Key

Per-DRF Variables

DRF_Portal_State (Indexed by Portal System Number)		
PSN = 1	PSN = 2	PSN = 3
DRF_Gateway_State	DRF_Gateway_State	DRF_Gateway_State
DRF_Aggregator_State	DRF_Aggregator_State	DRF_Aggregator_State

Maintain history of the state associated with unacknowledged sequence numbers for Home Portal System Number (and, if center of chain, for other Portal Systems).

Structures


Portal_State_Summary (Indexed by Portal System Number)		
PSN = 1	PSN = 2	PSN = 3
Gateway_Sequence_ Number	Gateway_Sequence_ Number	Gateway_Sequence_ Number
Aggregator_Sequence _Number	Aggregator_Sequence _Number	Aggregator_Sequence _Number

Per-IPP Variables

- Neighbor_Portal_State_Summary
- Neighbor_Gateway_State
- Neighbor_Aggregator_State
(stored from Home_xxx_TLV in RxDRCPDU)

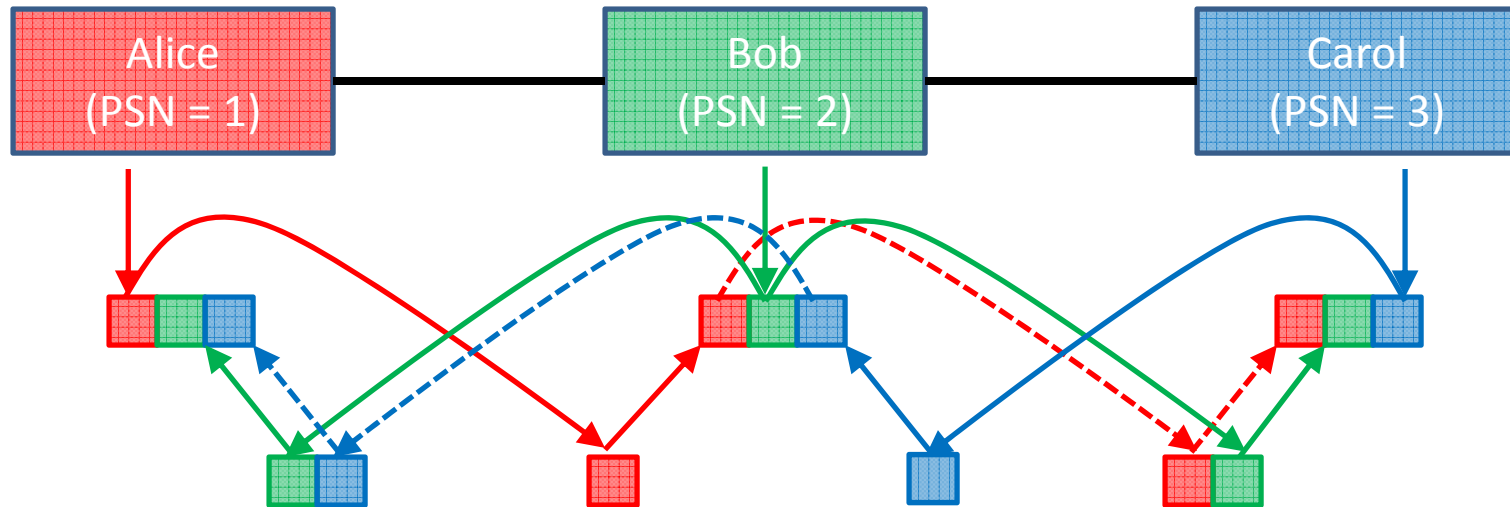
- Remote_Neighbor_Portal_State_Summary
- Remote_Neighbor_Gateway_State
- Remote_Neighbor_Aggregator_State
(stored from Other_IPP_Neighbor_xxx_TLV in RxDRCPDU)

Distributing Portal State (Three Portal System Chain)





- Home_Gateway_State_TLV and Home_Aggregator_State_TLV sent in DRCPDUs, and stored in Neighbor_Gateway_State and Neighbor_Aggregator_State variables at the receiving IPP. Copied to DRF_Portal_State variable.
- Other_IPP_Neighbor_Gateway_State_TLV and Other_IPP_Neighbor_Aggregator_State_TLV sent in DRCPDUs from the center of a three Portal System chain, and stored in Remote_Neighbor_Gateway_State and Remote_Neighbor_Aggregator_State variables at the receiving IPP. Copied to DRF_Portal_State variable at the end of a three Portal System chain.

Distributing Portal State (Three Portal System Chain)




Acknowledging Portal State (Three Portal System Chain)

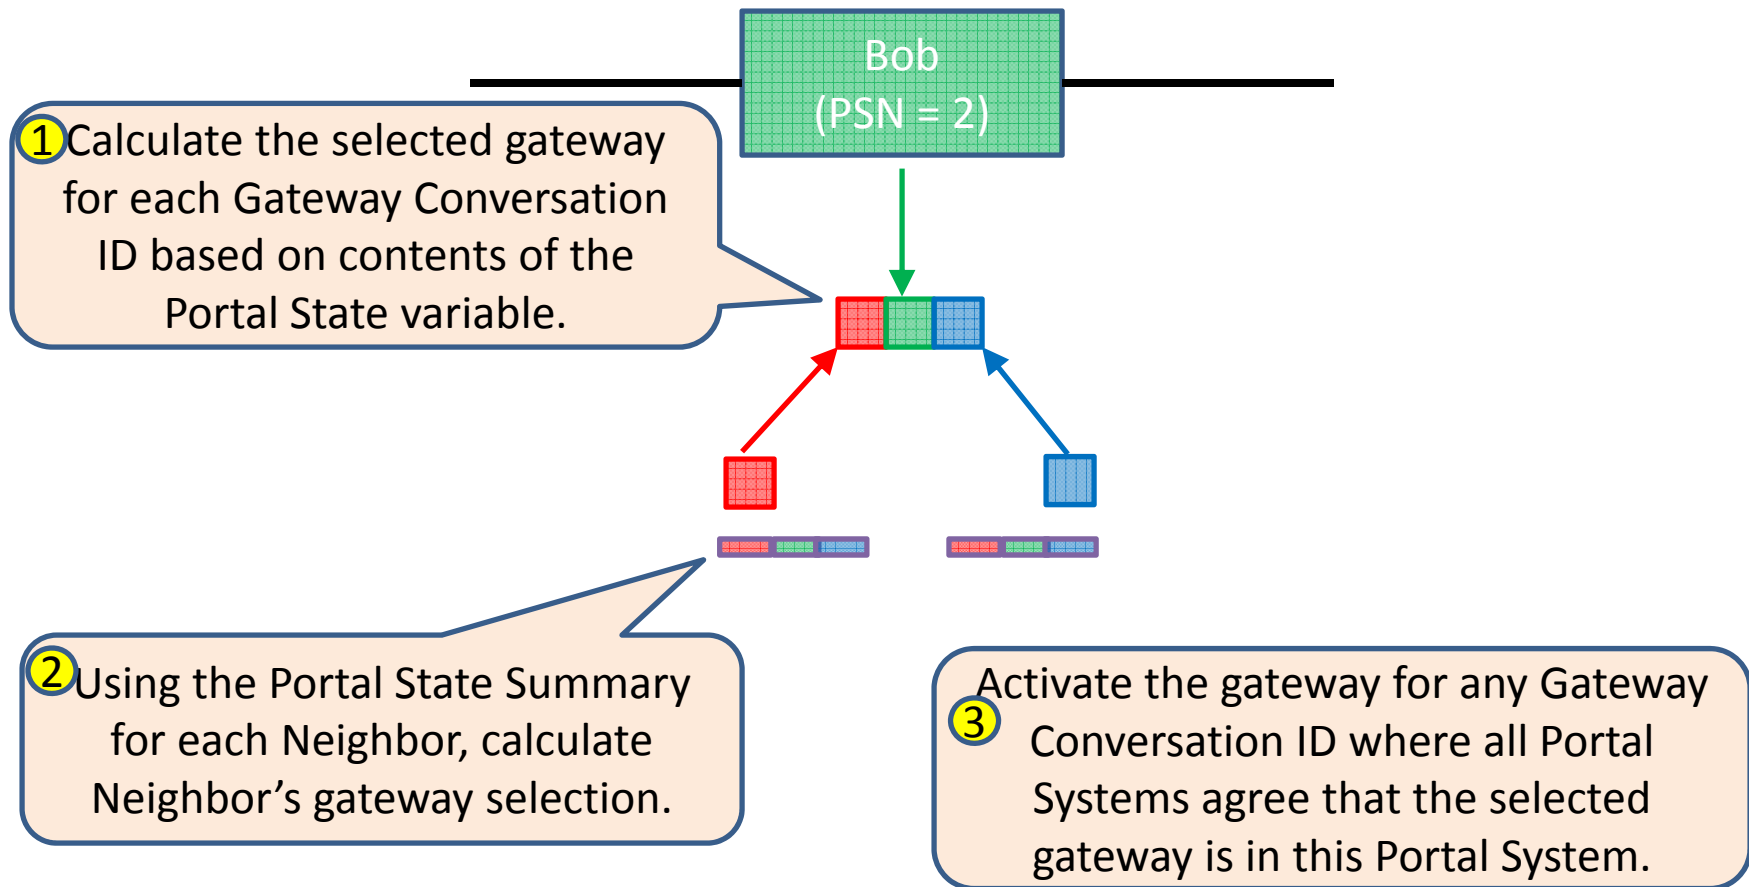


- • Home_Portal_State_Summary_TLV sent in all DRCPDUs, and stored in Neighbor_Portal_State_Summary variable at receiving IPP.
- - -→ • Other_IPP_Neighbor_Portal_State_Summary_TLV sent in DRCPDUs from the center of a three Portal System chain, and stored in Remote_Neighbor_Portal_State_Summary variable at receiving IPP.

Communicating Portal State (Three Portal System Chain)



Gateway Determination



Example gateway move

Initial condition: sync'd with all Gateway sequence numbers = 1.



Alice changes config (and sequence number).

2 1

Since "home" and "neighbor" gateway calculations give different results, Alice deactivates gateway for any CID that will move.

2 1
LACPDU

2 1
LACPDU

2 1

Bob deactivates gateway for any CID moving from Bob to Alice, and activates gateway for any CID moving from Alice to Bob.

Alice activates gateway for any CID moving from Bob to Alice.

2 1

Three Portal Systems?

- History
 - Support for three Portal Systems proposed as a way to enable hot-swap of one Portal System while always having at least two Portal Systems operational, so have system level resiliency even during maintenance intervals.
 - Support for three Portal Systems accepted as a “stretch” objective for the project: Include in the standard if it is not too complex.
- Feedback
 - Informal discussions with Carrier Ethernet Service Providers at the MEF Forum conclude that they would probably not use the three Portal System option, and would be satisfied having a single Portal Service operational while the other was undergoing maintenance.
 - “Too complex” is a subjective judgement. At this point it still appears that it is possible to get the three Portal System option to work*, however verifying correct operation (both in the standard and in testing an implementation) requires significantly(!!!) more work than just two Portal Systems.

Three Portal Systems? (cont.)

- More Feedback:
 - *Mick has done some analysis with a preliminary conclusion that correctly forwarding data frames between three Portal Systems may require labelling the frames as Up/Down frames and/or labelling with an association to a specific logical IPL (so a logical ring can be supported on a physical chain).
 - Rationale for continuing development of the three Portal System option is based on “momentum” rather than a strong desire for the feature. Informal discussions at 802.1 meetings conclude that interest in the feature ranges from “don’t care” to “absolutely don’t want it”, but not “absolutely want it”.

Proposal:

Eliminate the three Portal System option.

Mick's LACP proposal

<http://www.ieee802.org/1/files/public/docs2018/ax-seaman-lacp-suggestions-0118-02.pdf>

- LACP State Machine Overview
 - Informative diagram of state machine inter-dependencies
- Receive Machine
 - Eliminates “port_moved” variable.
 - Combines LACP_DISABLED and DEFAULTED states
- Mux Machine
 - Reflects Mux enhancements that were incorporated in draft 0.1 and updated in task group ballot comment resolution
- Transmit Machine
 - Combines function of old Periodic and Transmit State Machines
 - Restarts periodic timer after any LACPDU transmission