



# DRNI Port Selection

**Information flow for proper selection for sending frames to Gateway, Function, and IPL ports.**

**Norman Finn**


**Ver. 01**

# Reference

- This slide deck is available at:  
<http://www.ieee802.org/1/files/public/docs2012/ax-rev-nfinn-port-selection-0912-v01.pdf>
- Draft 0.3 of P802.1AX-REV is available at:  
<http://www.ieee802.org/1/files/private/ax-rev-drafts/d0/802-1AX-REV-d0-3.pdf>


# What's in P802.1AX-REV D0.3

# Topology restricts us to four cases:



- “I” do not know the state of the link between the other two Portal Systems.
- In case 1, I know that only two devices are configured.
- Case 2 is error condition for a triangle, normal for star.
- Case 3 may be a normal condition or an error condition.

# Up traffic and down traffic




- One **Gateway** per Portal System
- Any number of **Function ports** per Portal System
- One or more **Intra-Portal Link (IPLs)** per Portal System
- “Up” traffic is Function port to Gateway.
- “Down” traffic is Gateway to Function port.
- **Two** classification algorithms, one Up, one Down.
- **All of the problems are in distinguishing whether a frame received on an IPL is Up or Down traffic.**

# “Liberties”


- There are two kinds of liberties one can take when directing frames over IPLs:
- **Flood**: One can flood an up frame to one's own Gateway and IPL(s) and flood down frames to one's own Function port(s) and IPL(s).
- **Accept**: One can accept an up frame from a Function port or accept a down frame from a Gateway that, according to the other direction's rules, it should not be sent to.

# Flood and Accept liberties



- Every frame, no matter whether it is an up or down frame, can be classified and assigned a Gateway (or “none”) and a Function port (or “none”).
- **“Flood”** means the frame is sent to more ports than its destination indicates. (B sent frame left and right.)
- **“Accept”** means that the frame is accepted from a port that its source does not indicate. (C accepted the frame from port 6, when it’s a port 2 frame.)


# Too much liberty is chaos



- If System C takes the **Accept liberty** on **both** its Gateway and its Function ports, System B doesn't know which is the up frame for its Gateway, and which is the down frame for Function Port 4.




# Too much liberty is chaos



- If System A takes the **Flood liberty** and the **Accept liberty**, and floods a frame needlessly, then System B doesn't know whether to transmit the frame (a loop!) or discard it. (This could be an up frame from 1, 2, or 3 to A, or it could be a down frame from A to 4.)



# Just enough liberty



- The **Accept liberty** is allowed only for **up** frames.
- The **Flood liberty** is allowed only for **down** frames.
- The “Accept + Flood” problem is solved.

# Just enough liberty

- The **Accept liberty** is allowed only for **up** frames.
  - If a DR function receives a frame from a Gateway that it would not transmit on that Gateway, the DR function discards the frame.
  - A DR function is allowed to accept a frame from an unexpected Function port. (However, discarding that frame can ensure frame ordering.)
- The **Flood liberty** is allowed only for **down** frames.
  - A DR function transmits an Up frame only once, and only to the correct Gateway or IPL Port.
  - A DR function can flood a down frame to any combination of Function ports and/or IPL Ports. (However, this can be done only if the Function ports are able to discard unwanted frames.)

# A proposal for making decisions and changing them

# Who's in charge of this interface?

- Two separate configured items in a Portal for a Distributed Relay:
  - Either I'm in charge of Gateway selection, or Network is in charge of Gateway selection.
  - Either I'm in charge of Function port selection, or lower layers are in charge of Function port selection.
- Hopefully, the other side (Network or lower layers) are configured the opposite way.
- Hopefully, all of the DR functions in the Distributed Relay are configured the same way.

# Dynamic problem avoidance

- When Gateways are changing, it is essential that no data be forwarded in the direction towards the changing port (up).
  - Otherwise, differing ideas over where the data is going can lead to ambiguities in whether a frame is an up frame or a down frame, which means that frames can be looped or duplicated.

# Basic unit of information

- One can construct a vector of “which Portal System owns the entry/exit, if any,” indexed by Conversation ID.
- For a three-System Portal, this is a vector of four states (A, B, C, none) per Conversation ID.
- This basic unit of information can apply to which Portal System owns the Function port, or which owns the Gateway.



# Basic unit of information


- One side (Network or DR) tells the other what Gateway to use, the other side talks with its peers and, when ready, responds positively.
- One side (lower layers or DR) tells the other what Portal System's Function port to use, the other side talks with its peers and, when ready, responds positively.
- Because it requires additional state, and thus a larger vector, to accommodate multiple distinct Function ports, I would recommend that we say that there is only one Function port per DR function, and that multiple Functions below that one Function port have to be able to recognize their own traffic and discard other's traffic.

# The protocol

- On every interface (Function port, Gateway, or IPL), I present my current vector, and keep a copy of your current vector.
- This information is kept/transmitted separately for the two directions Up and Down. (These two may or may not have separate algorithms for mapping frames to Conversation IDs.)
- As long as I have a disagreement between me and the other side of the port for any port for a particular direction/Conversation ID, then I discard all frames on that direction/Conversation ID.
- If the other guy changes, I relay his changes, but do not respond with an agreement until I get agreement from the other ports.

# The one issue

# Issue with flooding



- If we allow three Systems in a Portal, and we allow flooding, System B must know whether the dashed **A-C IPL** is up or down.
  - If it's up, B **must not** relay flood data from C to A.
  - If it's down, B **must** relay flood data from C to A.
- In my opinion, conveying this knowledge amounts to a routing protocol that has convergence time. So, if flooding is allowed, I think three-System Portals should be disallowed, and vice-versa.
- I don't know for sure that this is not also a control-plane issue.