



LLDP TLVs required for P802.1au

Simplifying the Congestion Notification Domain interchange via LLDP

Version 2

Norman Finn

nfinn@cisco.com

References

- This presentation is available at:
<http://www.ieee802.org/1/files/public/docs2009/au-nfinn-LLDP-TLVs-0109-v01.pdf>

Defending CN Domains

- Receive defenses

A bridge prevents non-CN frames from entering a CN-enabled network on a CN priority by remapping all frames coming from a non-CN-enabled system to a non-CN priority, so that uncontrolled data streams do not cause uncontrollable congestion.

A station can prevent non-CN frames from entering a CN-enabled input queue because its functions that process these frames cannot handle non-CN-tagged frames.

- Transmit defenses

A station or bridge does not transmit CN-tagged frames from a queue on a CN priority if the receiving system is not CN-enabled, so that congestion-controlled streams can be sent to a system that is not CN-aware, and so that CN-tagged frames are not remapped to another priority by the receiver's defenses.

CN Domain protection events

1. NoMap

- The bridge or station turns off its remapping of priorities for ingress frames.

2. SndTags

- The bridge or station turns off its stripping of CN-tags on output.
- The station enables transmission of CN-tagged frames.

LLDP TLV **bits**, one bit per priority

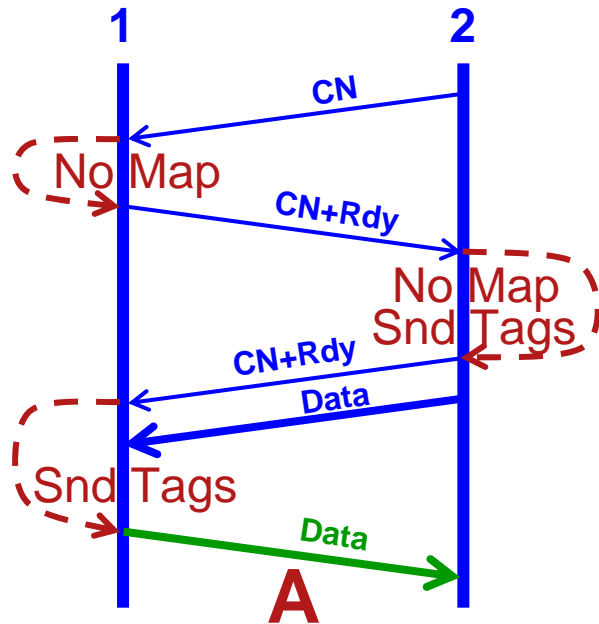
1. **CN[n]**

- This priority is a CN-priority.

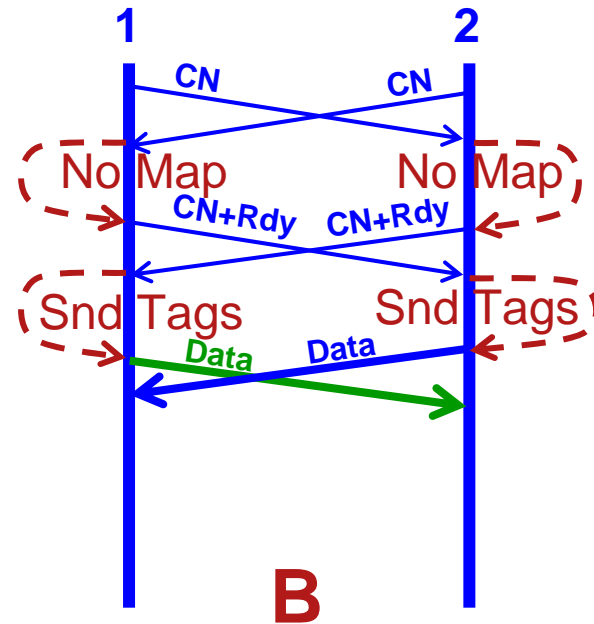
2. **Rdy[n]**

- I have turned off my priority remapping defenses.

Per-priority handshake: Both have receive defenses



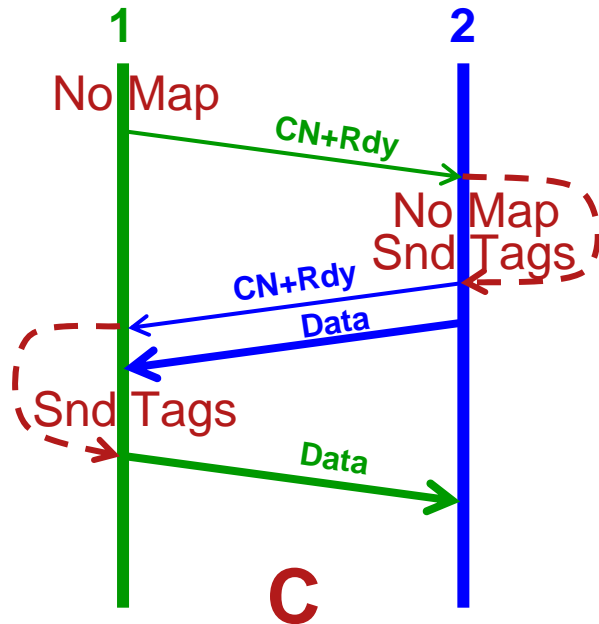
A
Device 2 first



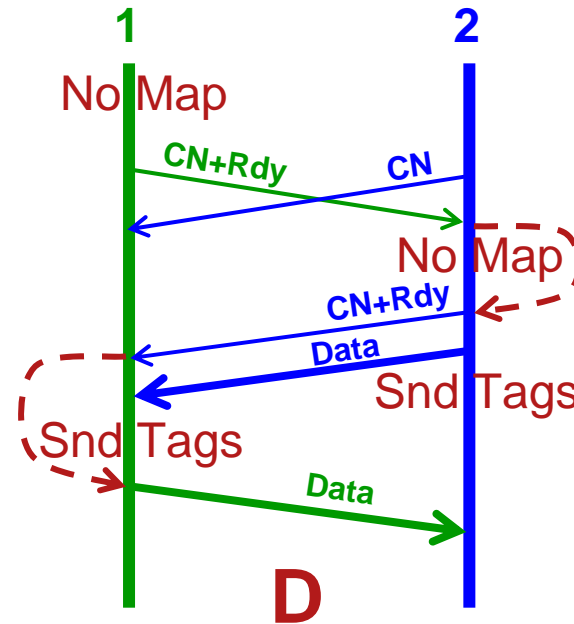
B
Neither first

- Each turns off its defenses (and turns on Rdy) after seeing “CN” in LLDP.
- Each enables tag sending after seeing the other’s Rdy bit.

Per-priority handshake: #1 has no receive defenses



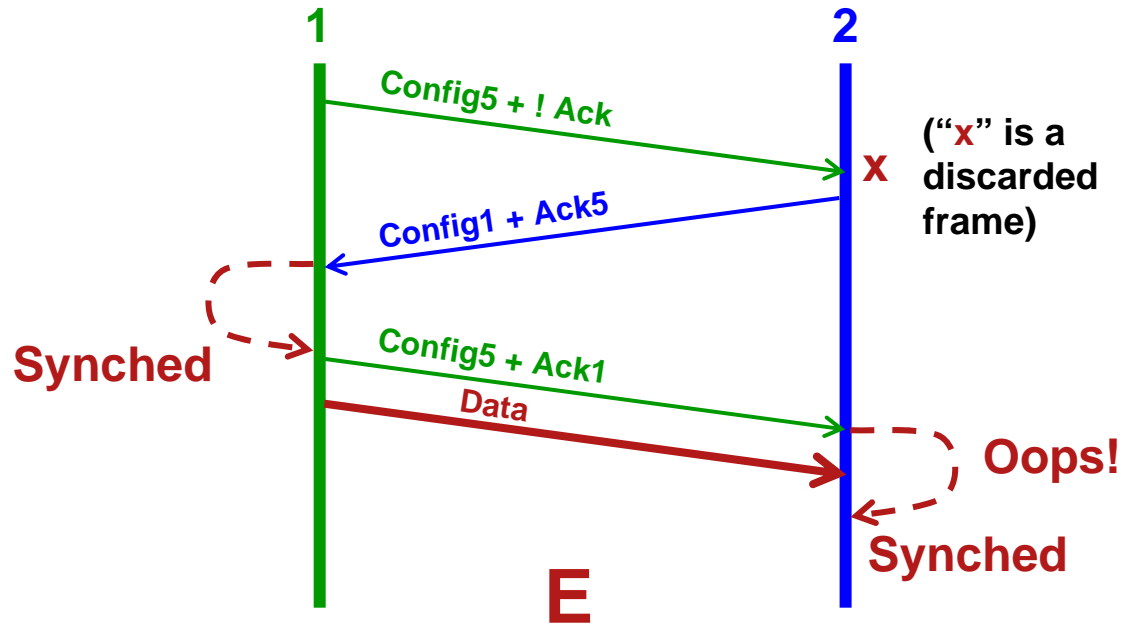
C
Device 1 first



D
Neither first
(or 2 first)

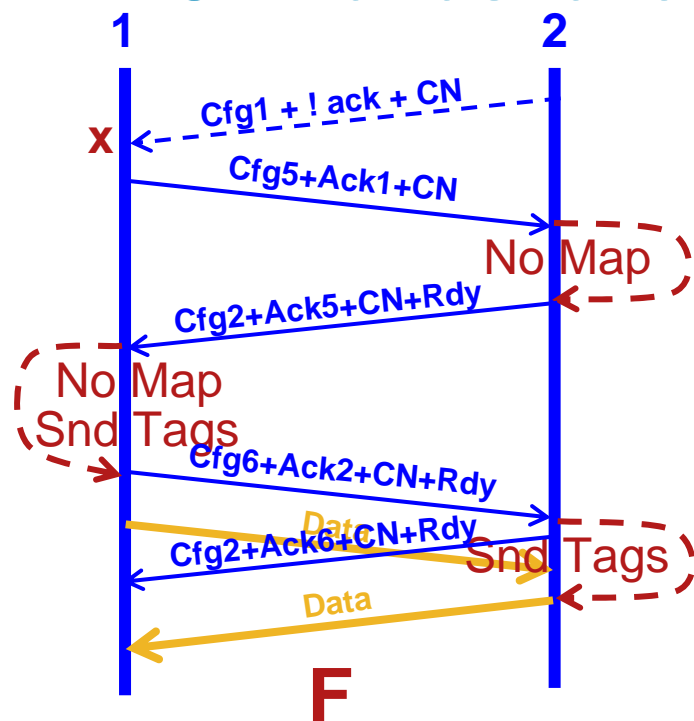
- #1 is always Rdy; it has no defenses to turn on.
- Neither can send tags until it sees the other's Rdy.
- Device #2's "CN" bit alone is of no interest to Device 1.

DCBX ACK handshake: race condition

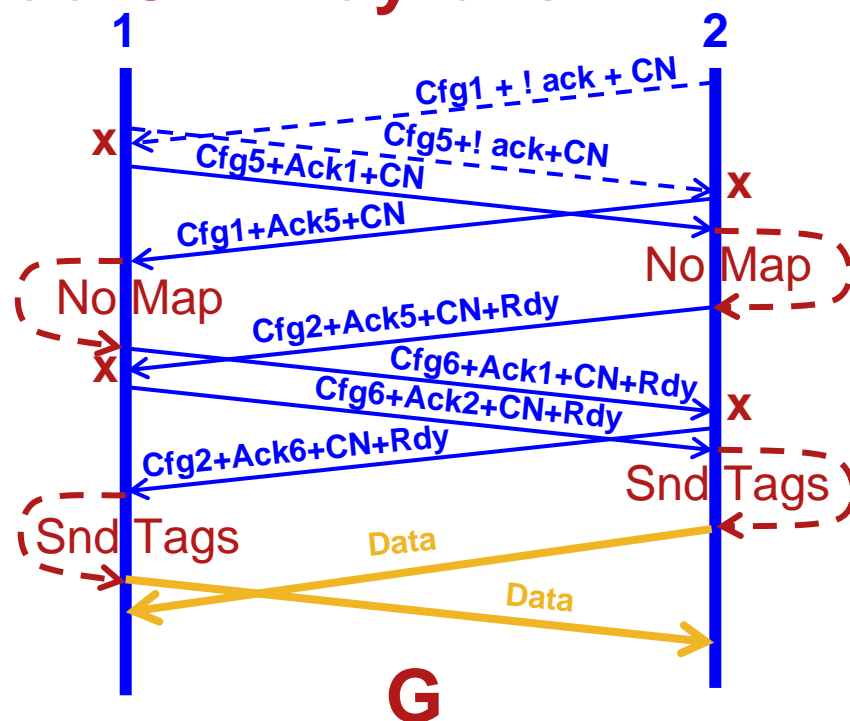


- If we just use the sequence numbers (a, b) to denote “turn on feature (turn off defenses and turn on tag sending),” then there is a gap during which device 1 sends before device 2 is ready to receive.

DCBX ACK handshake: Add CN+Rdy bits



F
Device 2 first



G
Neither first

- If we make the per-priority CN and Rdy bits a part of the DCBX feature payload, more frames must be exchanged.
- Case **F** takes **5** frames (**4** if pre-synched); case **A** took **3**.
- Case **G** takes **8** frames (**6** if pre-synched); case **B** took **4**.