Marvell. Moving Forward Faster

2.5G Link Fault Signaling

June 27, 2016

William Lo, Marvell

Purpose of Slides

- **Gives information on 802.3cb activity to help 802.3bz group make a decision on whether to make link fault signaling optional or keep it mandatory for 2.5GBASE-T**

- **Current rationale for making link fault signaling optional is based on the notion that the link fault signaling is not compatible with an extender SERDES using legacy 1000BASE-X PCS. This is not the case given the work of 802.3cb.**

Use of Sequence Ordered Set for Link Fault Signaling

- Clause 46.3.4 defines the Link fault signaling
  - Local Fault
  - Remote Fault
  - Link Interruption
- 802.3bz modified Clause 46 to apply to 2.5Gb/s and 5Gb/s speeds
- 10GBASE-T uses link fault signaling
  - Link Interruption indication is needed to stop MAC from transmitting when fast retrain is active
- 2.5GBASE-T inherits the 10GBASE-T behavior

Usage Model

- 2.5GBASE-T implementations most likely will not expose the XGMII.
- Most likely it will either be attached via a extender SERDES (i.e. 2.5GBASE-X) or directly integrated together with the MAC
- Extender will need to be able to pass sequence ordered set to pass link fault signals
- Extender will need to ensure start of packet is on 4 byte boundary as the 2.5GBASE-T PCS expects this

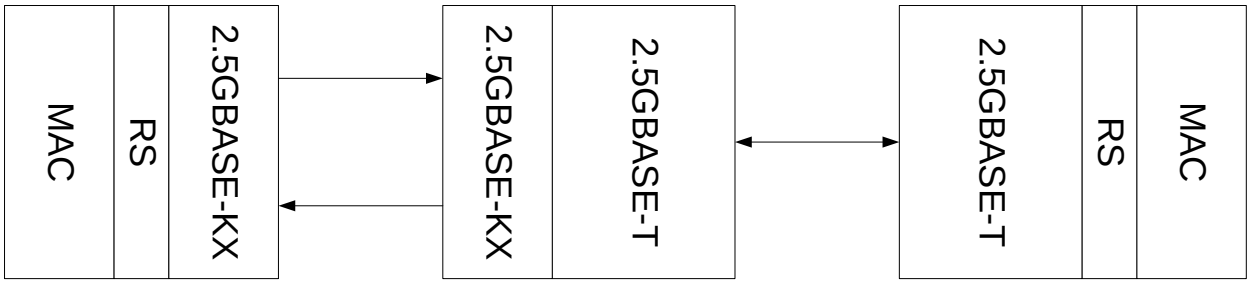

802.3cb Has a Solution to Send Sequence Ordered Set

- **2.5GBASE-KX starts with 1000BASE-X PCS and adds:**
  - **1 byte to 4 byte alignment needed to service XGMII**
  - **Passing Sequence Ordered Sets**
- **Implementation as shown does not require much incremental logic**

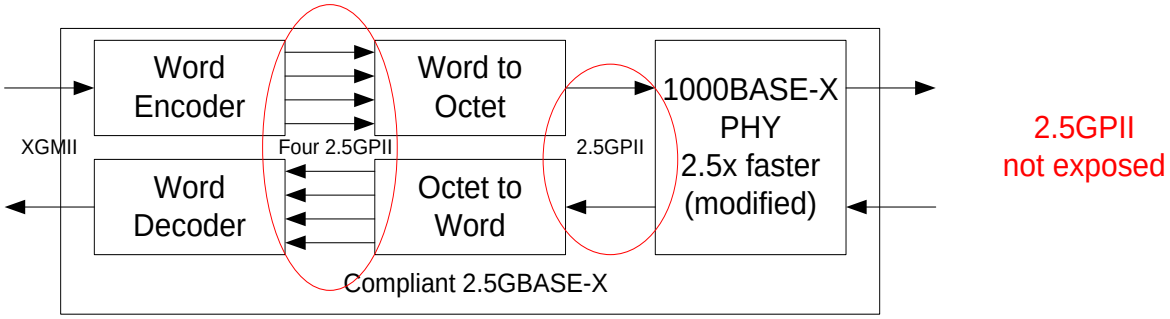

- **See Lo_3cb_01a_0116.pdf for theory of operation**

Sequence Ordered Set in 2.5GBASE-X

- **Sequence ordered set - |Q| = /K28.5/W/K28.5/W/K28.5/W/K28.5/W/**
- **To a 2.5x speedup 1000BASE-X PCS the |Q| looks exactly like 4 /K28.5/D/ idle symbols if |W| is not any of the following:**

Function	Data Code	Octet	7	6	5	4	3	2	1	0
LPI	D6.5	A6	1	0	1	0	0	1	1	0
LPI	D26.4	9A	1	0	0	1	1	0	1	0
Config	D21.5	B5	1	0	1	1	0	1	0	1
Config	D2.2	42	0	1	0	0	0	0	1	0

- **What this means is a legacy 2.5x speedup 1000BASE-X PCS can connect to a 2.5GBASE-X PCS and not break. The 1000BASE-X PCS simply sees idles instead of |Q|**

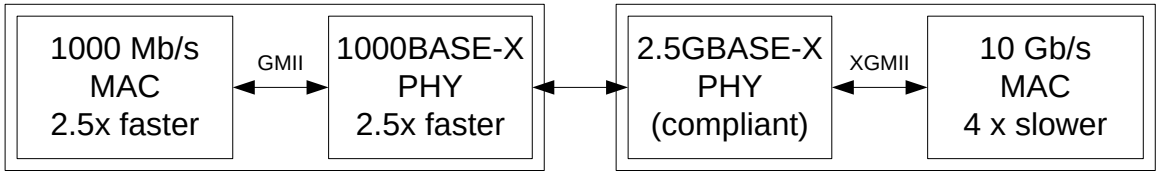

4 Permutations of Extenders with 2.5GBASE-T PHY

- **2.5GBASE-X to 2.5GBASE-X**
  - No problem. Link Faults Passed
  - 802.3cb saves the day
- **Legacy to 2.5GBASE-X**
  - Legacy sees idles even if 2.5GBASE-X sends link faults (see previous page)
  - No need to ask 2.5GBASE-T to turn off link faults
- **Legacy to Legacy**
  - Shim layer to align to 4 byte boundary
  - Shim layer can mask link faults to idles.
  - Rather than asking 2.5GBASE-T to behave differently, let shim layer block it to compensate for legacy limitations
- **2.5GBASE-X to legacy**
  - 2.5GBASE-X never sees link fault since legacy cannot sent it
  - 2.5GBASE-X never outputs link fault since it never has to respond to legacy link fault

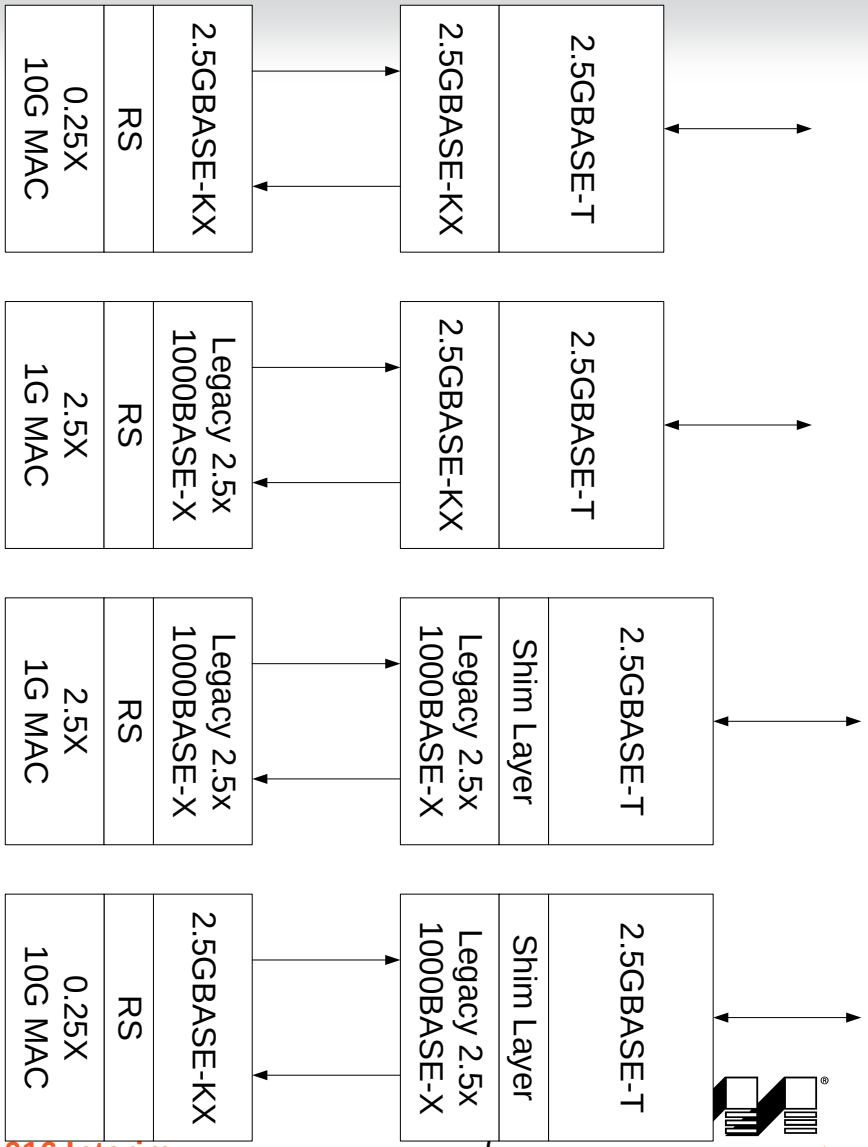

MARVELL

Conclusion

- There is no need to make 2.5GBASE-T link fault signaling optional to accommodate legacy 2.5x speed 1000BASE-X PCS
- 802.3cb already took care of the case where legacy MAC/PCS can work with a 802.3cb 2.5GBASE-X PCS
- If 2.5x speed 1000BASE-X legacy PCS is the extender interface in a 2.5GBASE-T PHY a shim layer is needed anyway and take care of blocking the link fault signaling outside the 2.5GBASE-T.

- Recommendation is to keep link fault signaling mandatory in 2.5GBASE-T to keep it consistent with 5GBASE-T, 10GBASE-T, 25GBASE-T, and 40GBASE-T rather than make an exception.

THANK YOU