IEEE 802.3bz
Making link fault signaling optional

Arthur Marris – Cadence

27th June 2016

Introduction

- These slides are updated from the slides presented to the ad-hoc on the 20th June
- I have modified them to only request making the link fault state machine optional for 2.5G.
- I have also expanded them to address some of the feedback I received during the ad-hoc call and afterwards

What is the problem?

- 802.3bz requires full implementation of the Clause 46 RS
- Clause 46 requires the implementation of link fault signalling
- The Clause 36 PCS (8B10B) used for SGMII does not support link fault signalling
- There are existing implementations that do not support the link-fault state machine
- There are existing implementations that should be able to connect to 2.5GBASE-T PHYs using the Clause 36 PCS

What needs to use link fault signalling?

- The fast retrain function sends the “link interruption” sequence ordered set to the MAC’s RS to cause it to defer transmission.
  - Fast retrain is important for 10G data rates but is less so for 2.5G.
- It can be used by BASE-T PHYs to recover without doing full auto-negotiation if there are problems during LPI

What is the solution?

- Make link fault signalling optional for 2.5G data rates.
- Add an extra sentence to the end of the paragraph on page 59, line 13, “For 2.5 Gb/s data rates implementation of link fault signaling as described in 46.3.4 is optional.”
- Bring subclause 46.3.4 into 802.3bz and change the last sentence from:
  - “The RS shall implement the link fault signaling state diagram (see Figure 46–11).”
- To:
  - “The RS shall implement the link fault signaling state diagram (see Figure 46–11) for data rates of 5 Gb/s and above. For 2.5 Gb/s data rates implementation of the link fault signaling state diagram is optional.”

What are the objections to making it optional?

- 2.5GBASE-T can use the link fault state machine to recover without doing auto-negotiation if there are problems during LPI
  - This is true, although when the state machine starts sending remote fault, data will be lost and the host system will be seeing a fault condition being reported from the link fault state machine
- It will be difficult to configure
  - Configuration can be done through MDIO
- You will fail compliance testing if you use speeded up SGMII
  - Seeing as the only interface specified is XGMII, you cannot fail compliance testing if XGMII is not exposed
- 802.3cb allows you to use SGMII if you use a shim layer in the PHY
  - This is true but I am not aware of this being previously discussed in 802.3bz

What is this PHY shim layer?

- Clause 46 introduces three extra requirements not supported at 1G rates:
  - Aligning start-of-packet on 4-byte boundaries
  - Encoding sequence ordered sets
  - Implementing the link fault state machine to respond to local and remote faults
- The 2.5GBASE-T PCS requires start-of-packet aligned on 4-byte boundaries so when using speeded up SGMII a shim layer is required in the PHY transmit path to implement deficit idle count to do the alignment
- If the link fault state machine is mandatory then the shim layer also needs to implement the link fault state machine if a legacy SGMII type extender interface is used
- Using a PHY shim layer of some kind is necessary if a legacy SGMII-type extender is used. This shim layer is more complex if link fault signalling is required.

Conclusion

- Requiring all future 2.5G implementations to support link fault signalling is quite a big deal
- It is up to the ballot resolution committee to decide whether to make it optional or not
- These slides are intended to increase understanding of the technical issues involved