

Addressing comment #71

Adee Ran, Cisco

Background

- Although PHY 400GXS is not mandatory, I assume typical 400BASE-ZR implementation is within a module connected with 400GAUI-n C2M.
- Comment #71 noted that the 400GBASE-ZR PHY and the PHY 400GXS are specified in a way that requires the AUI-C2M to always be active.
 - With Local Fault ordered sets sent over the AUI-C2M when the PHY is not receiving data correctly.
 - This is different from behavior of common modules (that do not include a PCS). The comment asked to highlight this difference.
- Meanwhile, P802.3df D2.0 comment resolution resulted in a significant change to link fault signaling over the 800GMII Extender
 - Allowing disabling of the module output as an optional behavior.

Should we allow squelch of module output in 802.3cw too?

- Disabling the output on the host interface when there is no optical signal is a likely module behavior
 - Specified as “Auto-squelch” in existing form-factor SFF documents and CMIS.
- Without squelching, the RS cannot distinguish between local fault sent by the link partner, and local fault generated by the PHY.
- Alignment with 802.3df (and future 802.3dj) would be beneficial.
- Relatively few simple changes are required (though they are spread over multiple clauses).
- Proposal in next slides should be implemented with editorial license.

Relevant changes in P802.3df D2.1 – PHY 800GXS

Expected changes in D2.1, not published yet, subject to change

- PHY_XS:IS_SIGNAL.indication was added in 171.3.2
- PCS_Status added to service interface below in 171.3.3
 - An additional signal, PCS_status, is provided in the receive direction (see 172.1.4.1)
- Added Link fault signaling subclause

The PHY 800GXS provides signal status information to the PMA using the PHY_XS:IS_SIGNAL.indication(SIGNAL_OK) primitive (see Figure 171–2). The SIGNAL_OK parameter is set to OK when PCS_status (see 171.3.3) is true and FAIL when PCS_status is false.




Figure 171–2—Functional block diagram for the PHY 800GXS

171.5 Link fault signaling in the receive direction

Link fault signaling generated by the local PHY or by the link partner is sent toward the Reconciliation Sublayer as ordered sets through the 800GMII Extender.

The PHY 800GXS provides signal status from the local PHY to the client PMA using the PHY_XS:IS_SIGNAL.indication(SIGNAL_OK) primitive (see 171.3.2 and Figure 171–2).

Proposed change: Apply corresponding changes for PHY 400GXS in Clause 118 with editorial license
No figures need to be updated

Relevant changes in P802.3df D2.1 – 800GBASE-R PCS

- PCS_Status was added to the service interface

- Figure 172–2 updated


Expected changes in D2.1, not published yet, subject to change

172.1.4.1 PCS service interface

The PCS service interface allows the 800GBASE-R PCS to transfer information to and from a PCS client. The PCS client is either the Reconciliation Sublayer or the PHY 800GXS sublayer.

When the client is the Reconciliation Sublayer, the PCS Service Interface is the Media Independent Interface (800GMII) in Clause 170.

When the client is the PHY 800GXS sublayer, the PCS Service Interface is the 800GMII in Clause 170 with an additional signal, PCS_status, in the receive direction. The PCS_status signal indicates the state of the PCS_status variable (see 119.2.6.2.2).



Proposed change: Apply corresponding changes for the 400GBASE-ZR PCS in Clause 155 with editorial license
Update Figure 155–3 adding PCS_Status output (aside from the 400GMII)

Relevant changes in P802.3df D2.1 – 800GBASE-R PMA

Expected changes in D2.1, not published yet, subject to change

- Provision was added to signal the status indication to a client PMA through an AUI by disabling the output of one or more lanes.
 - The relevant case for the PMA above the PHY XS is the 8:32 PMA:

173.5.8.2 8:32 PMA signal status

In the transmit direction the 8:32 PMA provides signal status information to the PHY 800GXS using the PHY_XS:IS_SIGNAL.request(SIGNAL_OK) service interface primitive (see 173.3 and Figure 173–4). The SIGNAL_OK parameter is set to OK when data is being received on all 8 input lanes (PMA:IS_UNITDATA_0:7.request). Otherwise SIGNAL OK is set to FAIL.

In the receive direction the 8:32 PMA optionally provides signal status information to the client PMA by disabling one or more output lanes (PMA:IS_UNITDATA_0:7.indication) when the PHY_XS:IS_SIGNAL.indication SIGNAL_OK parameter (see 173.3 and Figure 173–4) is FAIL.

Proposed change: Apply corresponding changes for the PMA in Clause 120 with editorial license (generic PMA for 200GBASE-R and 400GBASE-R; applies for all 400G PMAs, but 200G may not be in scope)

Relevant changes in P802.3df D2.1 – 800GBASE-R PMA

- Change made to Figure 173-4 showing that the SIL can disable the output.
- Figure 120-5 is the corresponding diagram for the 400GBASE-R PMA.
 - It is a more complicated diagram.

Applying a similar change in Figure 120-5 would be nice to have but not absolutely required.
Editorial license provided.




Figure 173-4—8:32 PMA functional block diagram

Relevant changes in P802.3df D2.1 – 800GAUI-8 C2M

Expected changes in D2.1, not published yet, subject to change

- Specification of module output in disabled state added to Annex 120G
 - Covers 400GAUI-4
 - But 400GAUI-8 is likely to be used too

120G.3.2 Module output characteristics

Change Table 120G-3 as follows (unchanged rows not shown):

Table 120G–3—Module output characteristics at TP4

Parameter	Reference	Value	Units
...			
Differential peak-to-peak output voltage (max)	120G.5.1		
Short mode		600	mV
Long mode		845	mV
Disabled		35	mV
...			

Proposed change: Apply corresponding change for the 400GAUI-8 C2M in Annex 120E with editorial license

Additional change required in P802.3cw

- Delete the highlighted sentence in 120.3 (PMA service interface, in 802.3-2022)
 - Because the PHY XS can provide its status to a colocated PMA.
 - Implement with editorial license.

The PMA:IS_SIGNAL.indication primitive is generated through a set of Signal Indication Logic (SIL) that reports signal health based on receipt of the *inst*:IS_SIGNAL.indication from the sublayer below, data being received on all of the input lanes from the sublayer below, buffers filled (if necessary) to accommodate Skew Variation, and symbols being sent to the PMA client on all of the output lanes. When these conditions are met, the SIGNAL_OK parameter sent to the PMA client via the PMA:IS_SIGNAL.indication primitive has the value OK. Otherwise, the SIGNAL_OK primitive has the value FAIL. In the case where the sublayer below the PMA is a PHY XS the PMA does not receive a PHY_XS:IS_SIGNAL.indication as an input to the SIL.

That's all!

(With editorial license)