ILT path start-up in Physical Layers that include Extenders

Adee Ran, Cisco

Introduction

- This presentation includes a detailed proposal to enable a path start-up procedure that spans PHYs and xMII extenders.
- It is based on communicating ISL status across a PHY XS + xMII + PCS, similar to the way it is communicated between the two interfaces of a PMA.
- The proposal covers the technical details and assumes broad editorial license for implementation in the next draft.
  - Also, it does not go into detail about potential scope issues.

Previously…

Cl **178B** *SC* **178B.5.3** *P* **789** *L* **44** # 421

Ran, Adee — Cisco Systems

Comment Type **TR** *Comment Status* **R** *(Common) ILT extender*

The text about training xMII extenders does not address the communication of the status variables isl_ready and remote_rts between interfaces (PMD to AUI and vice versa) when there is a PHY XS and PCS between them.
Ideally, this communication should be the same as the one defined in 178B.14.2.1 using adjacent_signal_ok, but the case of an extender is not covered by NOTE that describes what "adjacent" is.

Since this behavior is specific to PHYs attached to extenders, it should be specified in this subclause, preferably with a diagram.

SuggestedRemedy

Add a NOTE in 178B.5.3 stating that, for the purpose of adjacent_signal_ok, the adjacent interface of a PMD in a PHY attached to an xMII extender is the service interface of the PHY XS; and the adjacent interface of the AUI component above the PHY XS is the service interface of the PMD.
Add a figure to illustrate the communication of adjacent_signal_ok between the PMD and the AUI (across the PCS and PHY XS, and possibly other sublayers).

Response *Response Status* **C**

REJECT.

The CRG reviewed slides 24 to 28 in the following contribution:
https://www.ieee802.org/3/dj/public/25_07/brown_3dj_03a_2507.pdf

Straw poll TF-1 (below) shows strong consensus to define startup signaling that extends RS to RS.

However, the proposed solution does not provide sufficient detail to implement at this time. For instance, it is missing details for exchanging signals across the PCS service interface.

A detailed contribution on this subject is encouraged.

Straw poll #TF-1 (directional)
I support the direction of extending path start-up signaling (as proposed in D2.0 comment #421) from Reconciliation sublayer to Reconciliation sublayer.
Yes: 23
No: 1
Abstain: 20

A simple proposal to enable ILT signaling to encompass PHY XS+PCS, as if they were a PMA

There was clear consensus on the direction, but a more formal proposal was requested.

What we want to achieve

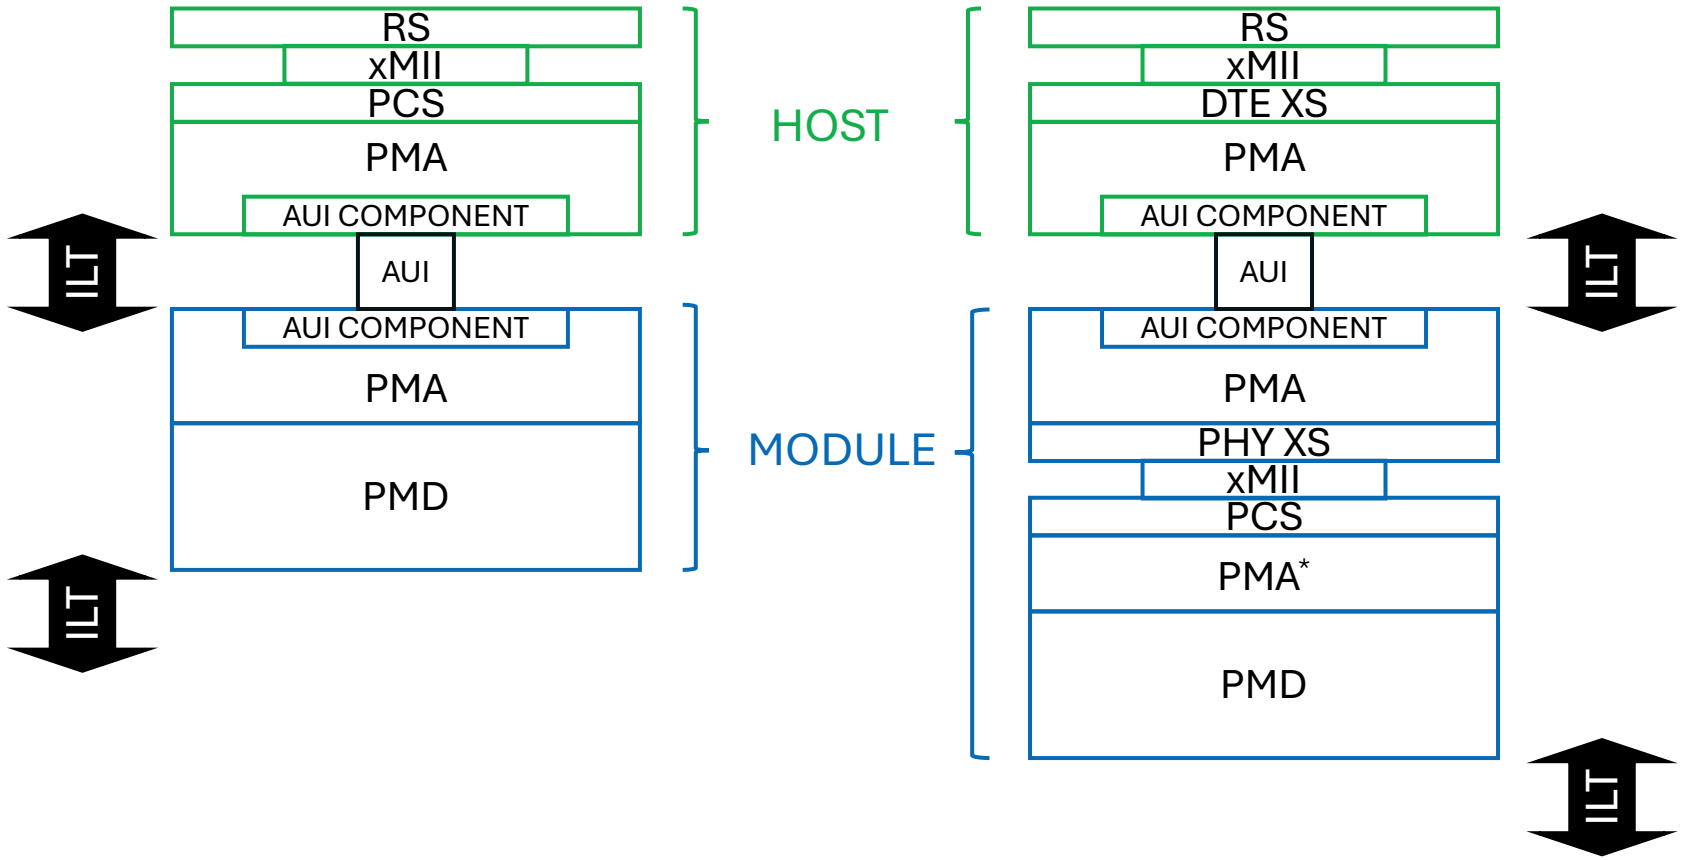

The implementation of ILT in the host and module AUI components, and the path start-up procedure, should be the same regardless of whether the module includes a PHY XS or not.

* Coherent PHYs have other sublayers, but the idea is the same (this presentation is not about adding ILT to coherent)

Illustration of the previous proposal

Proposed response – Option 1

NOTE – For the purpose of adjacent_signal_ok, the adjacent interface of a PMD in a PHY attached to an xMII extender is the service interface of the PHY XS; and the adjacent interface of the AUI component above the PHY XS is the service interface of the PMD.

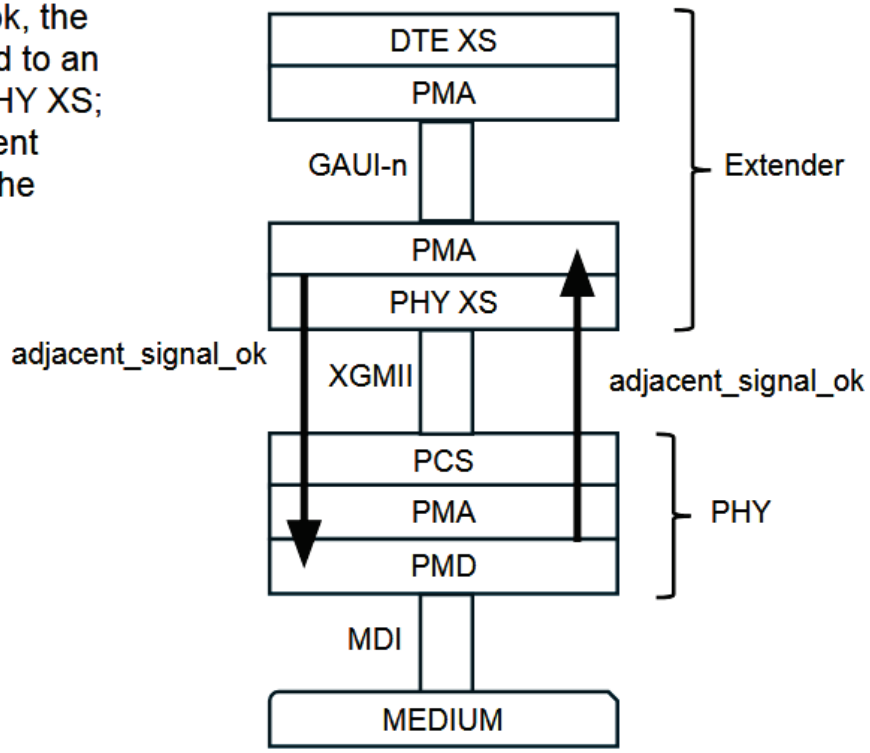

The idea was to "bypass" the intermediate sublayers, as if the PMD was directly below the bottom PMA of the extender.

However, this would be a violation of the sublayer separation, and an unclean specification.

Source: brown_3dj_03a_2507, slide 27

Path start-up essential primitives: IS_SIGNAL

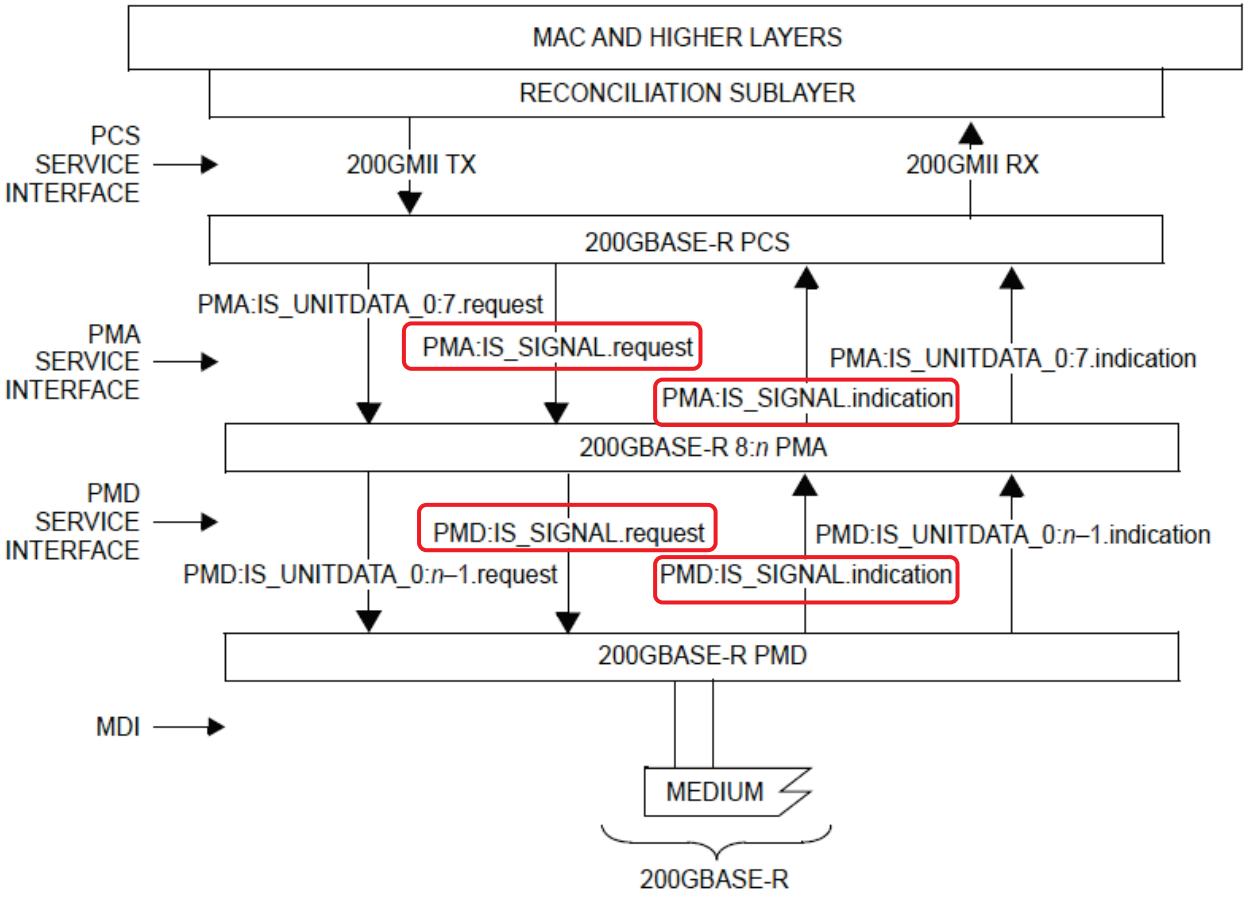

Figure 116-2

The **IS_SIGNAL.indication** and **IS_SIGNAL.request** primitives carry the ISL status information (including RTS) between sublayers.

But there are no equivalent signals in the PCS service interface (xMII).

How a PMA communicates SIGNAL_OK between its interfaces

Table 176–5—PMA:IS_SIGNAL.indication(SIGNAL_OK) generation

inst:IS_SIGNAL.indication[a] SIGNAL_OK	align_status_mux[b] or all_locked_demux[c]	PMA:IS_SIGNAL.indication SIGNAL_OK
OK	true	OK
OK	false	READY
READY	don't care	READY
IN_PROGRESS	don't care	IN_PROGRESS
FAIL	don't care	FAIL

[a] From the sublayer below the PMA.
[b] For n:m PMAs (see 176.4.4.2.1).
[c] For m:n PMAs (see 176.4.4.2.1).

Table 176–6—inst:IS_SIGNAL.request(SIGNAL_OK) generation

PMA:IS_SIGNAL.request[a] SIGNAL_OK	align_status_mux[b] or all_locked_demux[c]	*inst*:IS_SIGNAL.request[d] SIGNAL_OK
OK	true	OK
OK	false	READY
READY	don't care	READY
IN_PROGRESS	don't care	IN_PROGRESS
FAIL	don't care	FAIL

[a] From the sublayer above the PMA.
[b] For m:n PMAs (see 176.4.4.2.1).
[c] For n:m PMAs (see 176.4.4.2.1).
[d] To the service interface below the PMA.

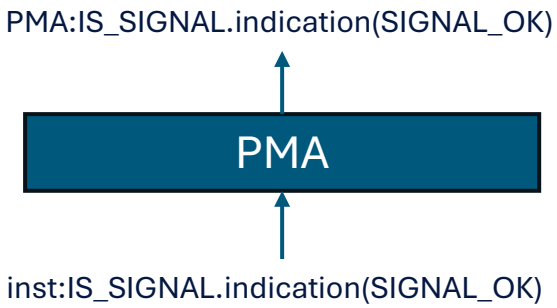

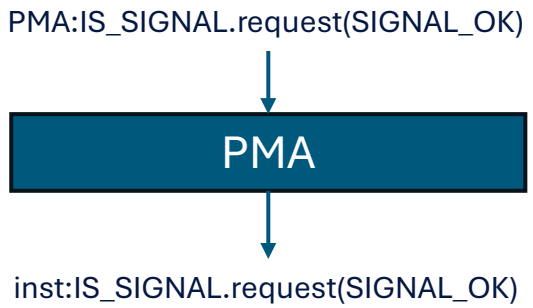

Note: there are no similar rules for the Clause 173 and Clause 120 PMAs

PCS service interface

Top of Figure 119-2
(In 802.3dj; this part is the same as in 802.3-2022)

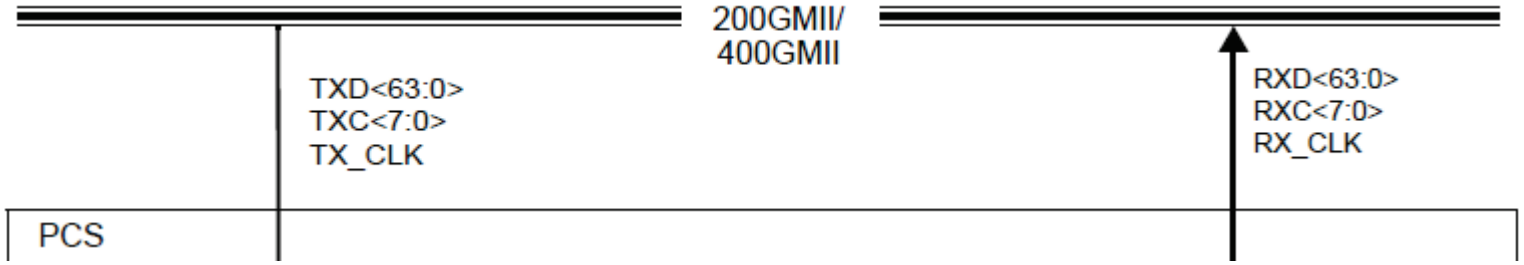

Top of Figure 172-2
(In 802.3dj; this part is the same as in 802.df)

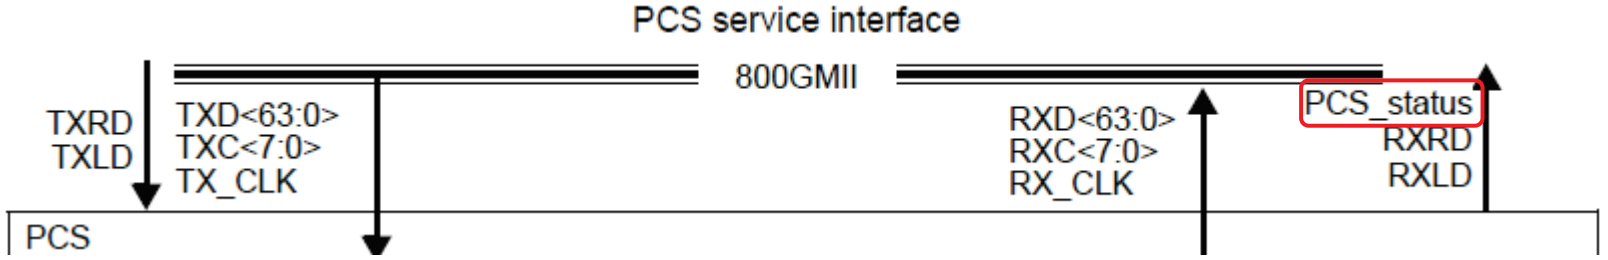

Top of Figure 175-2
(new in 802.3dj)

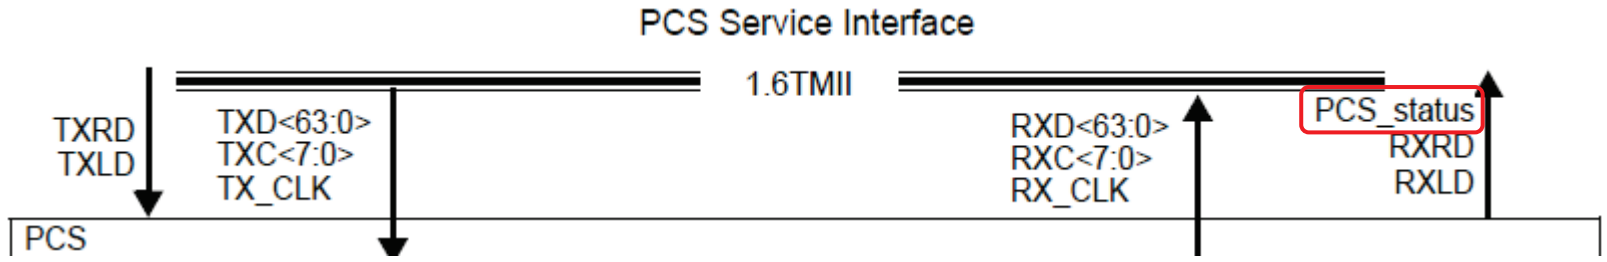

1.6TBASE-R PCS service interface

175.1.4.1 PCS service interface

The PCS service interface allows the 1.6TBASE-R PCS to transfer information to and from a PCS client as illustrated at the top of Figure 175–2.

When the client sublayer is the Reconciliation Sublayer, the PCS service interface is the 1.6 Tb/s Media Independent Interface (1.6TMII) defined in Clause 170.

When the client sublayer is the PHY 1.6TXS (see Figure 171–1), the PCS service interface is the 1.6TMII with additional signals TXRD, TXLD, RXRD, RXLD, and PCS_status. The TXRD and TXLD signals indicate the state of the remote degrade and local degrade, respectively, as detected by the PHY 1.6TXS in the transmit direction (see 171.3.3a). The RXRD signal indicates the state of the rx_rm_degraded variable (see 175.2.5.5) as detected by the PCS in the receive direction. The RXLD signal is the logical OR of the FEC_degraded_SER variable (see 175.2.5.3) and the rx_local_degraded variable (see 175.2.5.5) as detected by the PCS in the receive direction. The PCS_status signal indicates the state of the PCS_status variable (see 175.2.6.2.2).

PCS_status

Boolean variable that is true when align_status is true and is false otherwise.

Ideally **PCS_status** could convey the full status of the link, like SIGNAL_OK… However

- It is defined as a Boolean value, based on the PCS_status variable, which only considers align_status. It does have all possible values of SIGNAL_OK.
- Also, unlike SIGNAL_OK, it is defined only in the receive direction (from a PCS to a PHY XS).
- It has existing meaning that we should not break.

Interface below the PCS?

175.1.4.2 Service interface below the PCS

The service interface below the PCS supports the exchange of encoded data between the PCS and the sublayer below the PCS. It is an instance of the inter-sublayer service interface defined in 174.3 and is described in an abstract manner which does not imply any particular implementation.

Bottom of Figure 175-2 (new in 802.3dj)

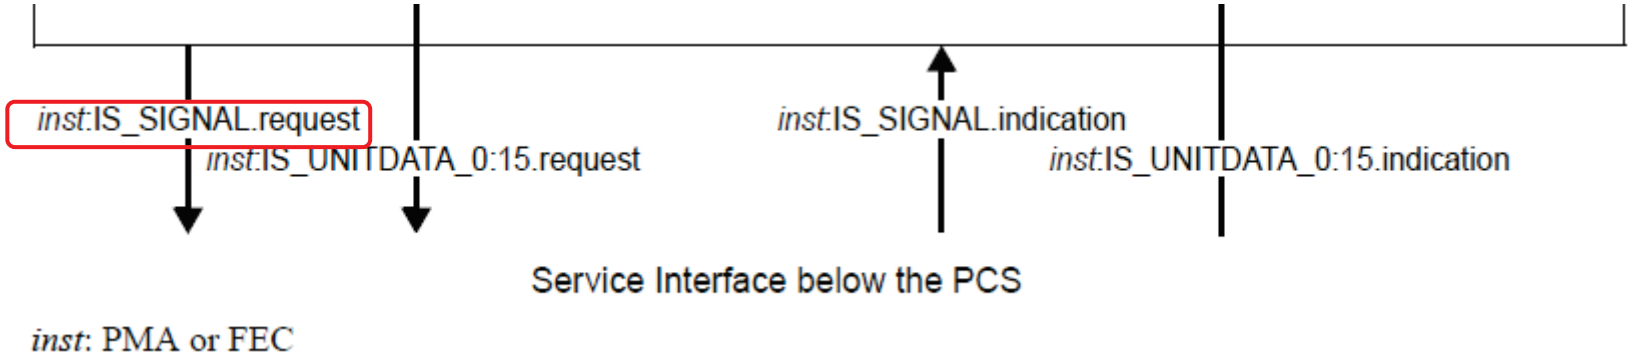

inst: PMA or FEC

Text in 175.2.1

Transmit data units are sent to the service interface via the *inst*:IS_UNITDATA_0:15.request primitive. The SIGNAL_OK parameter of the *inst*:IS_SIGNAL.request primitive is set to OK when the transmit data units are valid and is set to FAIL otherwise.

Similar diagrams and text exist in clauses 119 and 172.

General solution

- To enable passing the SIGNAL_OK across the PCS and PHY XS, the PCS service interface should be extended
  - Only when its client is a PHY XS
  - The PHY XS is an inverted PCS, so the interface below it will inherit the same definitions
- In the receive direction (from PCS to PHY XS) we have PCS_status, but it is insufficient
- Use a new signal, rx_link_status (with editorial license)
  - An enumerated value, logic similar to that of the link_status parameter in 119.6, but with addition of the READY state
  - The rules for generating PHY_XS:IS_SIGNAL.indication(SIGNAL_OK) should be modified to use rx_link_status, in alignment with Table 176-5
- Add a similar signal tx_link_status (with editorial license) in the transmit direction (from PHY XS to PCS)
  - Based on Table 176-6
  - The PCS rules for generating inst:IS_SIGNAL.request(SIGNAL_OK) should be modified to use tx_link_status, in alignment with Table 176-6
- If a sublayer does not generate the new signal (e.g., existing device), the receiving sublayer interpretation of the missing value is implementation dependent.

Definition of new signals

rx_link_status generation, based on Table 176-5

inst:IS_SIGNAL.indication[a] SIGNAL_OK	align_status	rx_link_status
OK	true	OK
OK	false	READY
READY	don't care	READY
IN_PROGRESS	don't care	IN_PROGRESS
FAIL	don't care	FAIL

From the sublayer below (PMA or other)

Use a similar definition (based on Table 176-6) for the signal tx_link_status.

Changes where?

- Clause 175 (1.6TBASE-R PCS)
- Clause 171 (PHY XS)*
  - Inherits the functionality of the PCS, but signal names and service interface direction are different
- Clause 172 (800GBASE-R PCS)*
  - PHY XS is also in clause 171
- Clause 119 (200GBASE-R and 400GBASE-R PCS)*
- Clause 118 (PHY XS)*
- No technical effect in Annex 178B, but change the text in 178B.5.1.2 (which is descriptive) to better align with the new concept that the path extends from MAC to MAC.

* The project scope should be addressed appropriately.

Summary

- A general proposal for communicating the ISL status information across the PCS and PHY XS (including their interface, the xMII).
  - Add new service interface signals tx_link_status (from XS to PCS) and rx_link_status (from PCS to XS), modeled after the SIGNAL_OK parameter of PMA:IS_SIGNAL.indication.
  - Generation of SIGNAL_OK service interface parameters in the PCS and in the XS is based on the new signals and rules similar to those of the PMA.
- Slides 11-13 should be implemented with editorial license.

That’s all

Questions?