Addressing possible 800G copper cable objective

Adee Ran, Cisco

Supporters

- Chris Cole, II-VI
- Itamar Levin, Intel
- Howard Heck, Intel
- Sam Kocsis, Amphenol
- Chris DiMinico, PHY-SI
- Ali Ghiasi, Ghiasi Quantum LLC

Goal of this presentation

- Lay out architecture options for supporting optics and copper ports, including both 100 and 200 Gb/s per lane electrical interfaces.
- Examine technical feasibility of dual-purpose (optics and copper) ports
  - Combining technical feasibility of AUIs, optics, and copper cables, including interoperability and implementation flexibility.
- We don’t need to select a solution at this point!
  - … but show that there are solutions we can choose from.

Background

- The market need for 800G-by-4 Ethernet over various optical media has been shown in multiple contributions and is established
  - Presentations addressed technical feasibility: kuschnerov_b400g_01_210503, mi_b400g_01a_210517, lam_b400g_01a_210720
  - 800G-by-4, 1.6T-by-8, and also 400G-by-2 and 200G-by-1 objectives have been adopted
    - Architecture is expected to be similar for PHYs with different number of lanes; similar signaling
- Objectives for Electrical interfaces (AUIs) for 200 Gb/s per lane have been adopted
  - Technical feasibility has been discussed (lyubomirsky_b400g_01_210329, healey_b400g_01a_210329, lu_b400g_01_210322) with techniques including PAM4 and higher order modulations
  - Related discussions about possible FEC enhancements with some architecture options (e.g. he_b400g_01_210426)
  - Commonality with optical signaling and existing architecture is preferable
- We also have objectives for 100 Gb/s per lane versions of 800G and 1.6T PHYs and AUIs
  - Objectives for 8x100 Gb/s copper cable and backplane have been adopted (kocsis_b400g_01a_0812)
  - These will likely be extensions of PHYs and AUIs being defined in 802.3ck

200G Copper cable

- 200 Gb/s per lane copper cable PHYs at 200, 400, 800, and 1600 Gb/s have recently been proposed objectives as based on market needs (tracy_b400g_01a_210729)
  - Technical feasibility was discussed in lu_b400g_01b_210729, noujeim_b400g_01_210517 – suggesting signaling methods other than PAM4.
  - Other approaches may enable using PAM4 over copper cables with reaches of 1.5 m or higher.
  - Technical discussion should continue in the task force.

Dual-purpose ports

- Having dual-purpose or universal ports (pluggable form factor supporting both copper cables and optical modules of various types) is highly desirable.
  - For short-reach links, this has been the norm for several generations, providing flexibility to users.
  - Expands market potential for both optics and copper.
- In 802.3ck (100 Gb/s per lane) we are getting to the limit of host loss for CR ports…
  - In switches, some ports may be unable to support CR without an external retimer.
  - Expect more pain with 200 Gb/s per lane!
  - Better start thinking sooner than later
- What paths do we have to enable universal ports?

Path #1: common signaling method

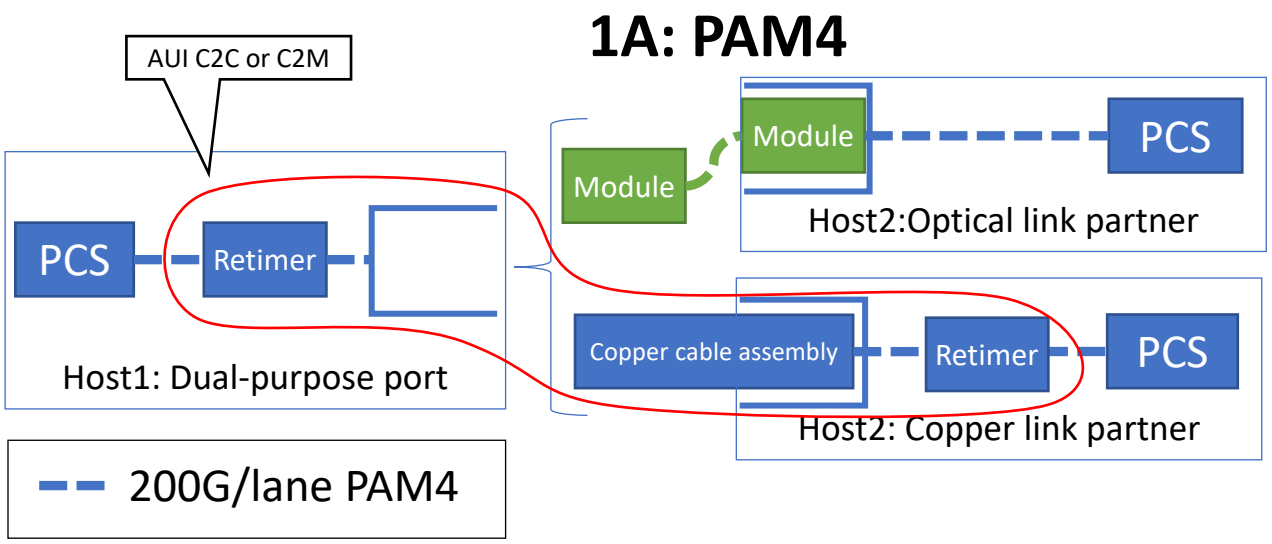

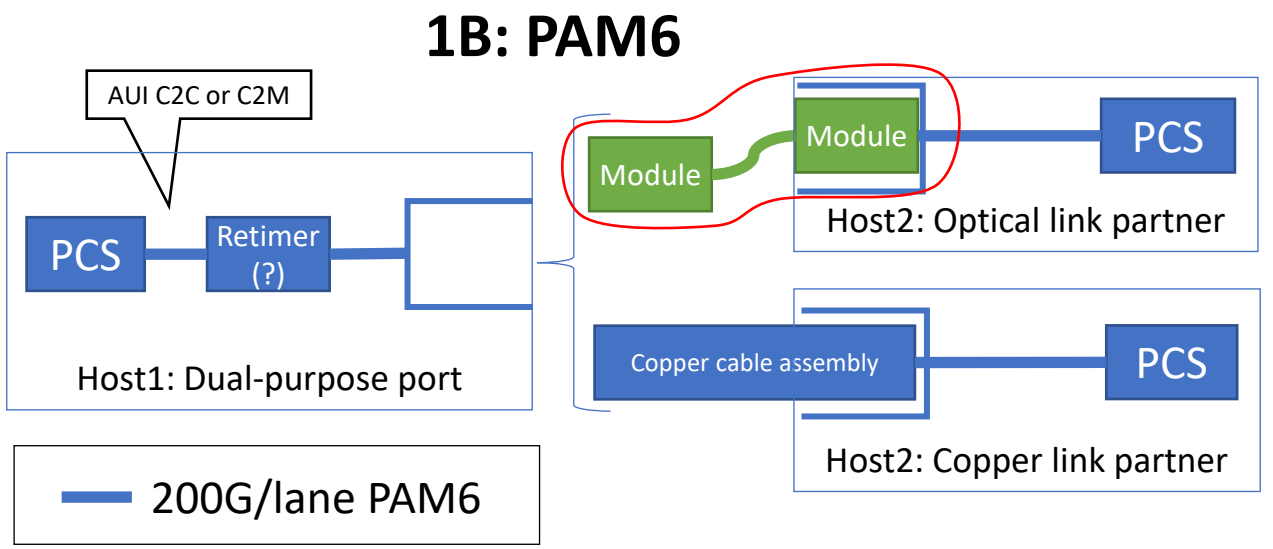

- Short-reach copper cable channel may be enabled by low-loss designs or using retimers

- Optical channel feasibility with PAM6 unlikely

Note: Further considerations for 200G/lane PMDs with 100G/lane host IO should be considered due to inclusion of necessary gearbox functionality but not shown in this presentation for simplicity

Path #2: different signaling for optics and copper, dual-purpose ports

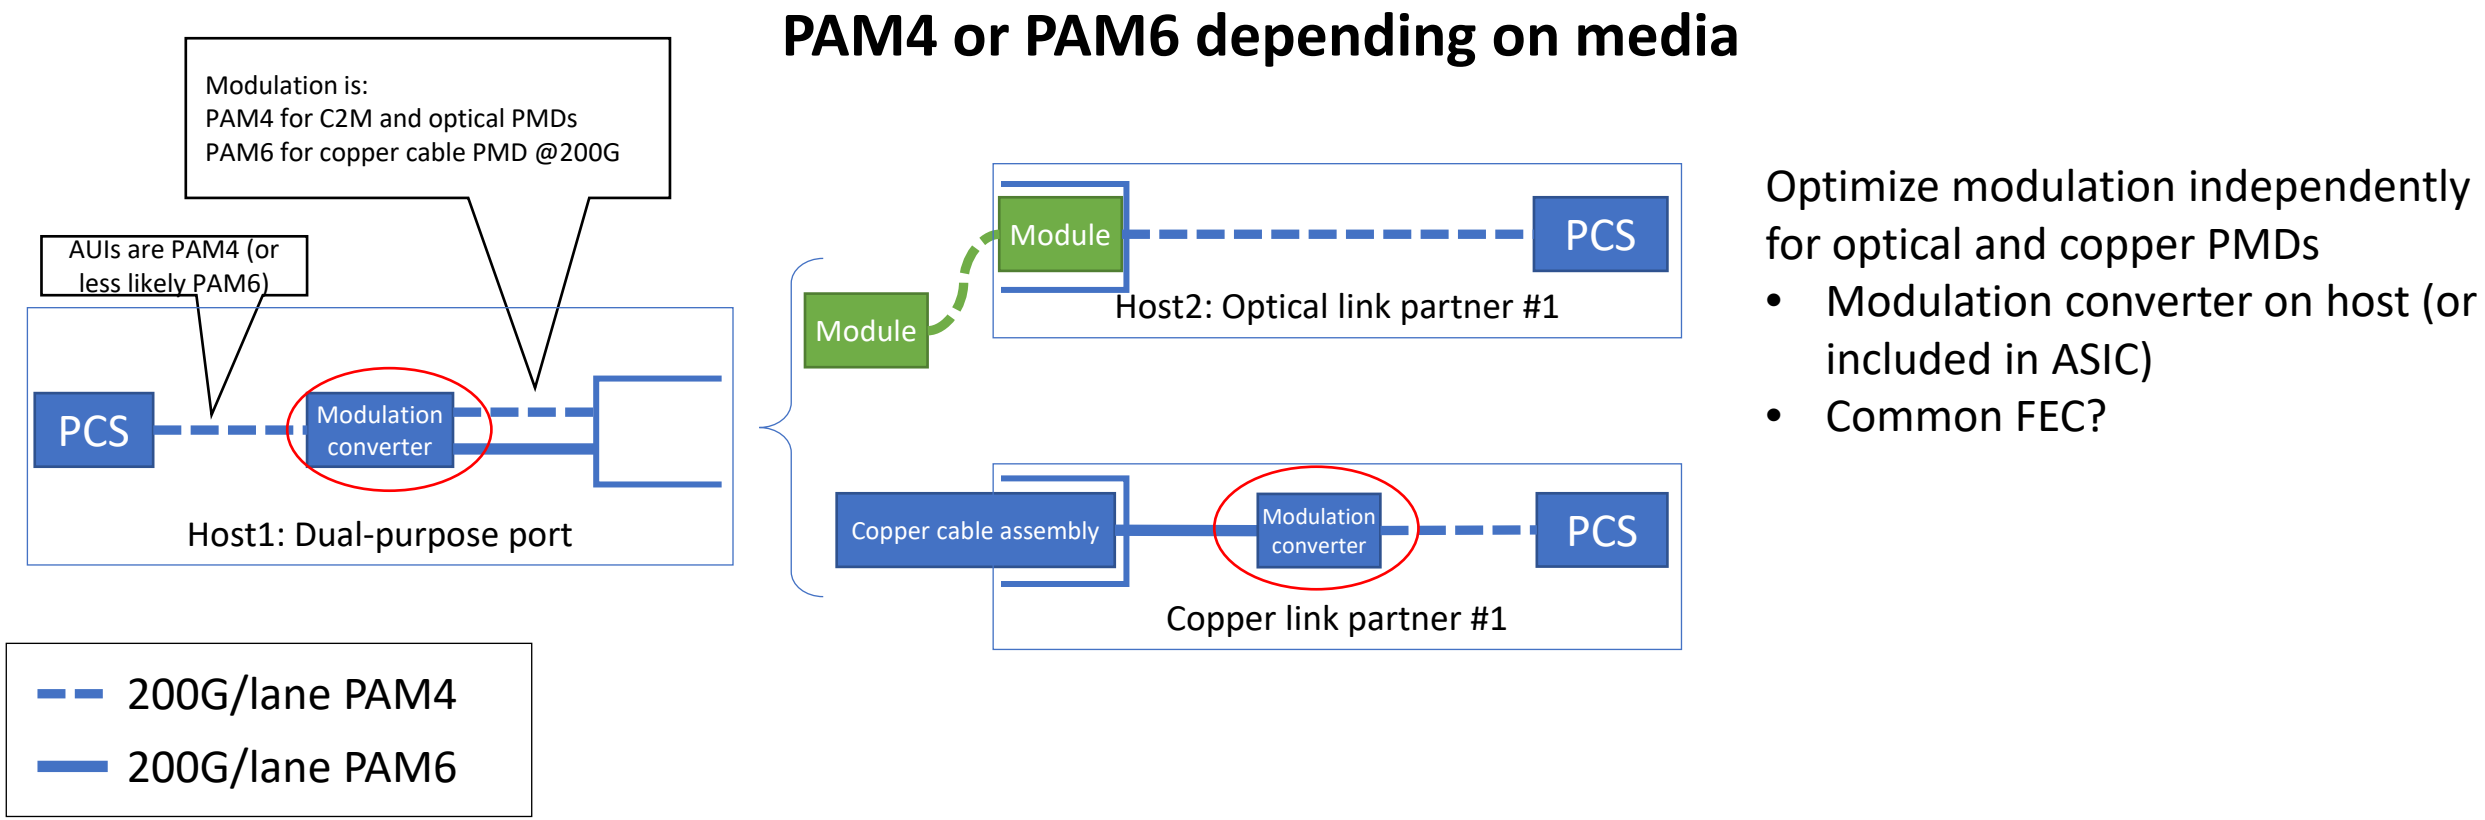

Optimize modulation independently for optical and copper PMDs

- Modulation converter on host (or included in ASIC)
- Common FEC?

What is a “modulation converter”?

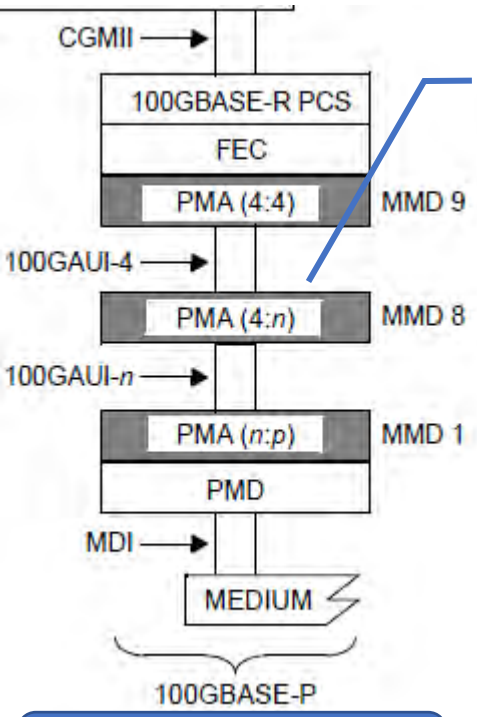

Excerpt from Figure 135–2 (802.3ck D2.1)

Bit mux*

- PMA with electrical interfaces on both ends
  - 1x200 Gb/s PAM4 on one end
  - 2x100 Gb/s PAM4 on the other
- Below the FEC
  - Assuming FEC is implemented in the same device as the PCS
- Synchronous 2:1 gearbox
- “Simple to implement”, protocol agnostic

Modulation converter

- Architecturally, also a PMA
  - 200 GB/s with PAM6 modulation on one end
  - 200 Gb/s or 2x100 Gb/s with PAM4 modulation on the other
- Maps symbols (or groups) to bits and vice versa
  - More logic, power and latency than a bit mux
  - Protocol-specific
- A device may need to function as a bit mux (PAM4 to PAM4) if different modulations are used for optics and copper

* Note: Not shown in the diagrams in this presentation but a bit mux would be needed for 200G/lane PMDs with 100G/lane host IO

Path #3: different signaling for optics and copper, dedicated ports

PAM4 or PAM6 depending on port

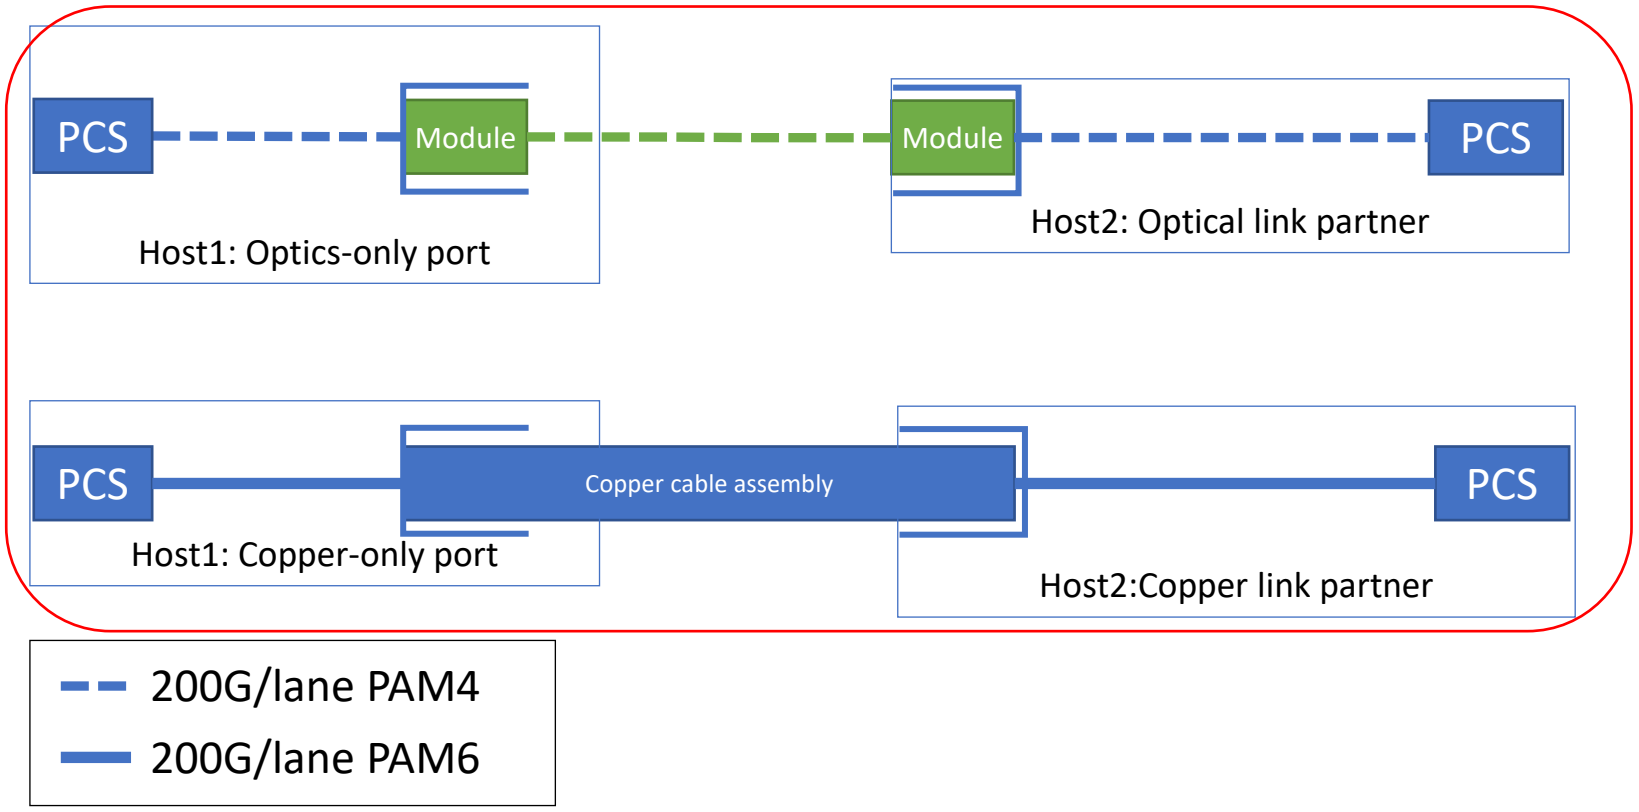

Optimize Ports for different PMDs.
Could be implemented using:

- Dedicated hardware switch ports for each PMD
- Dedicated Host ASIC (or modes)
  - multiple PCS/FEC blocks
  - Dual or complicated single SerDes designs

Path #4: Optical modules and active cables

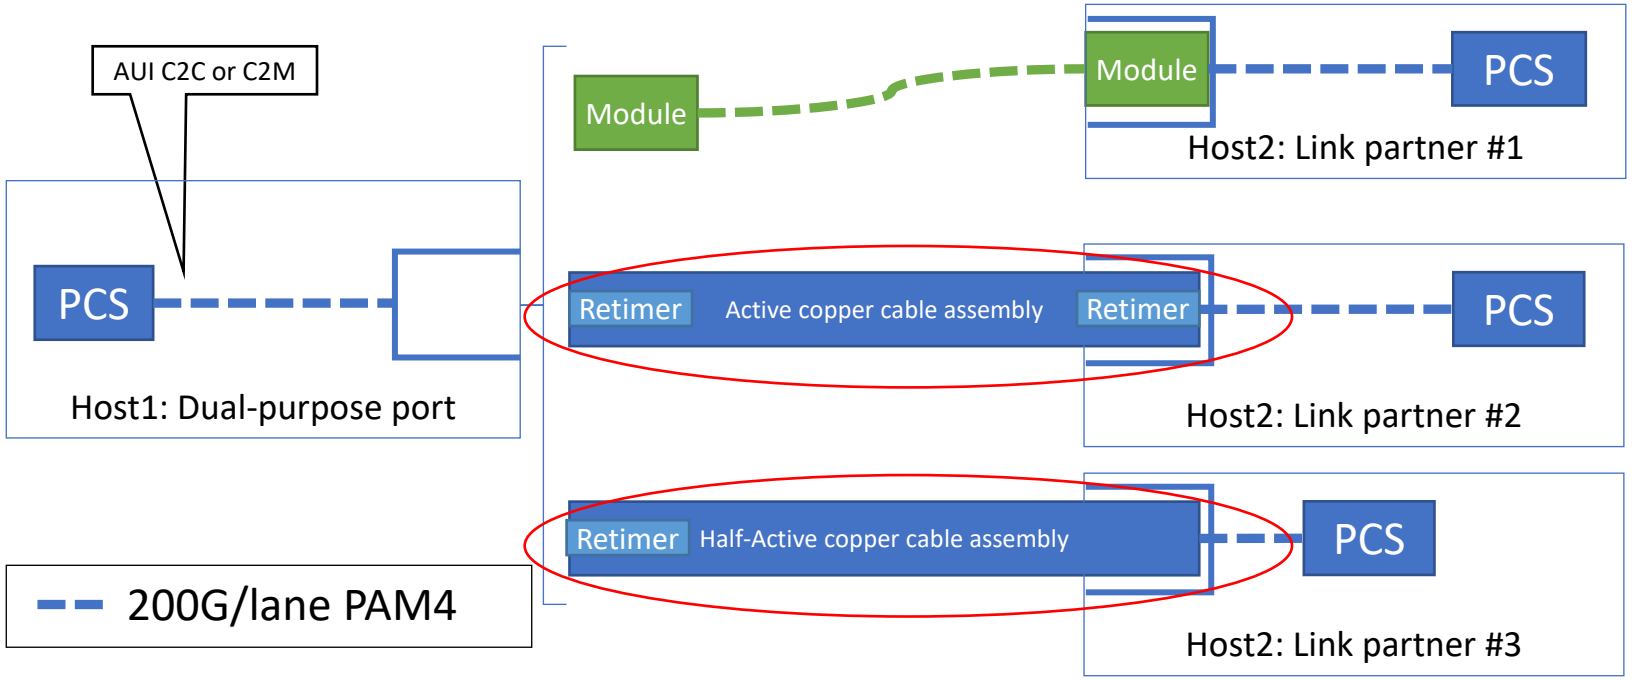

Copper Active cables allow the support of copper cables while architecture stays consistent with the optical PMD architecture

- Common host ASIC and system architecture
- In some cases, half-active cable (retimer on only one end) may work, reducing cost/power. Partitioning the channel's loss may allow common modulation with optics.

Architecture with active cables

- The host has a similar AUI-C2M interface to whatever is plugged in. Thus, all ports are dual-purpose.
  - FEC-encoded data over 100 Gb/s lanes can be bit-muxed together to form 200 Gb/s ports, regardless of medium
  - Simpler compatibility with 100G/lane ports and devices.
  - No need for modulation converters or retimers on the host to enable copper.
- No need to allocate loss budget for CR link.
  - BER is budgeted instead – as in optics.
  - Simpler design and analysis of each component.
  - Better confidence of system performance when each component is tested for compliance separately.
- Additional advantages
  - The channel is logically segmented. Each segment has lower loss, and the cable segment can internally use several signaling solutions (PAM4, PAM4SE, PAM6), and FEC schemes (segmented, concatenated, or end-to-end) transparently.
    - Using these advantages may increase reach, enable more use cases, expand the market potential, and drive down costs.
    - Note that passive cables can only use end-to-end FEC scheme; this may affect choice for optical PMDs.
  - Lower power consumption on the big ASICs compared to end-to-end links. Host can power an optical module anyway. Essentially, retimers are moved from the host to the cable.

Should 802.3 standardize active cables?

- Not considered necessary in the past.
  - Requirements can be implied from AUI-C2M (electrical) and optical module (BER) specifications
- However, explicit standard specification and nomenclature are important for customers.
  - Management as an example…
- As we have done in several other “obvious” PMD types, this would be a good service to the industry.

Summary

- Dual-purpose ports are technically feasible, but copper cable and optical interfaces should be designed jointly as a system solution.
- Architectures for supporting dual-purpose ports can be defined with either all-PAM4, all-PAM6, or dual-modulation schemes.
- Active cables can help in defining a clean architecture with easy interoperability and compatibility, such that all ports support both media types.
- Passive cables with shorter reach may also have wide market potential and be technically feasible.