Architecture: Configurations and use cases

Configurations and use cases ad-hoc

John D' Ambrosia & Hugh Barrass

with assistance from Pete Anslow, Mark Gustlin, Adam Healey, David Law, Gary Nichol, David Ofelt

Ad hoc participants

- Hugh Barrass, Cisco; John D'Ambrosia, Dell
- Pete Anslow, Ciena; Mark Gustlin, Xilinx; Adam Healey, Avago; David Law, HP; Gary Nichol, Cisco; David Ofelt, Juniper

- Hassan Ali, Texas Instruments
- Brad Booth, Microsoft
- Matt Brown, AppliedMicro
- Sue Bueti, IBM
- David Chen, NSN
- Piers Dawe, Mellanox
- Mike Dudek, QLogic
- Ali Ghiasi, Ghiasi Quantum
- Steve Gorshe, PMC-Sierra
- Adam Healey, Avago
- Paul Kolesar, Commscope
- Yonatan Malkiman, Mellanox
- Tom McDermott, Fujitsu
- Rich Mellitz, Intel
- Ray Nering, Cisco
- Tom Palkert, Luxtera
- Scott Powell, ClariPhy
- Rick Rabinovich, Alcatel-Lucent
- Adee Ran, Intel
- Salvatore Rotolo, ST Microelectronics
- Jeff Slavick, Avago
- Daniel Sparacin, Aurrion
- Peter Stassar, Huawei
- Nathan Tracy, TE
- Steve Trowbridge, Alcatel-Lucent
- Alexander Umnov, Fujitsu

General picture

- Many (most) host systems will include a C2C and a C2M interface

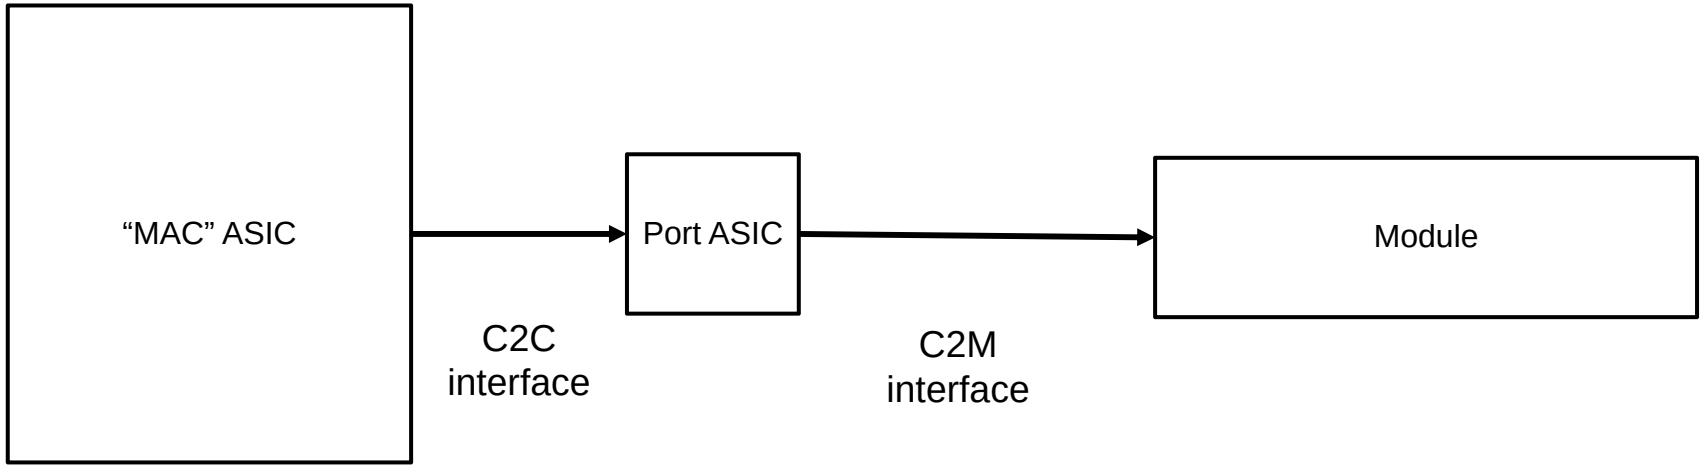

"MAC" ASIC includes MAC & a bunch of other stuff – often serves multiple ports

Port ASIC contains newer technology that could not be accommodated in current generation of "MAC" ASIC.

Module will be as simple as possible (with thermal/space challenges) – also accommodates new PHY components that weren't available when hosts system finalized.

General architecture

- Basic components from gustlin_bs_02_0514

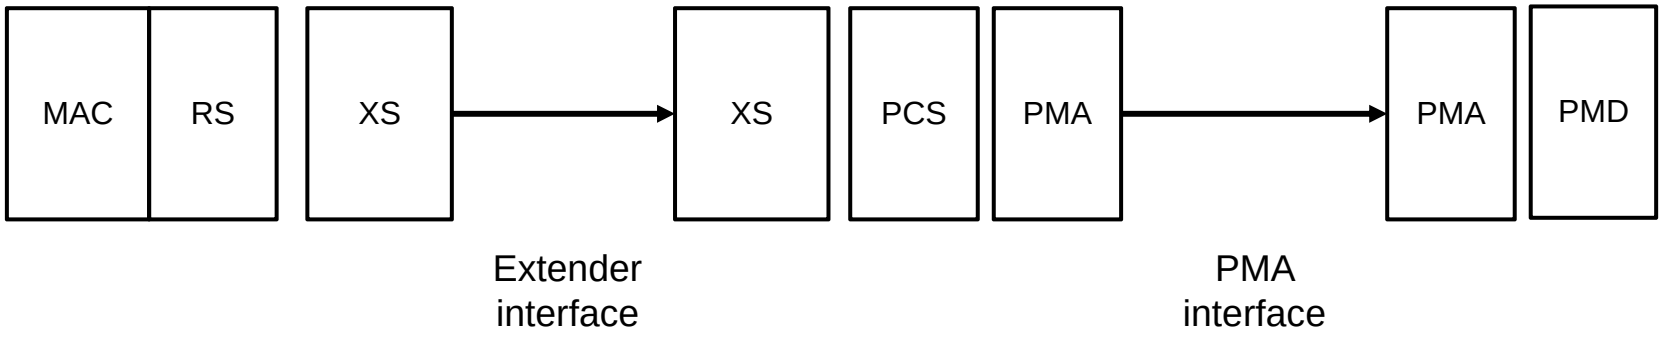

XS and extender interface is used to extend the xMII so that the raw data stream can be presented to the PCS

Extender is used in any situation where the coding or FEC requirements are different between the internal interface and the external PHY

PMA interface is used when there is no requirement to change the coding or FEC between the internal interface and the external PHY

Naming the interfaces

- CDAUI & CDXI

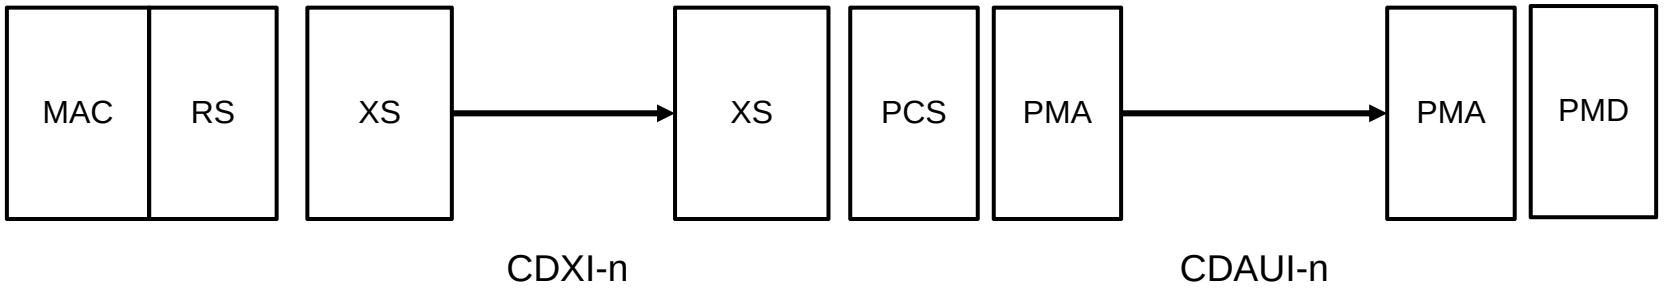

AUI – "Attachment Unit Interface" – fancy-speak for "the way you connect to a pluggy-thing"

Generally, if a host has both an extender interface AND a PMA interface, then the PMA interface will connect to a module – therefore PMA interface is an "AUI"

Unfortunately, there will also be cases where a host has 2 extender interfaces or 2 PMA interfaces – we just have to live with that…

From a practical perspective – the electrical specifications for both interfaces should be similar (so that a single chip can serve both functions), this will lead to use of "CDAUI" as shorthand term to refer to any instance of 400G physical interface (similar to XAUI in 10G).

Architectural definitions for the purists

CDAUI – 400G Attachment Unit Interface = physical instantiation of PMA interface

CDXI – 400G Extender Interface = physical instantiation of extender interface

Configs

- Capturing possible current/future configurations

Mapping host configuration examples to architectural diagrams – aim to build a comprehensive library of configurations

Capture host sublayer breakdown, with various module options – with variations at both ends of a link

Detailed proposals for coding/FEC/etc. can be measured against forward & backward compatible implementations

Current assumption is that multiple FEC will be achieved using segment-by-segment, note that implementations may be allowed to optimize this for pass-through cases

e.g. a 400GBASE-SR16 could be drawn like this…

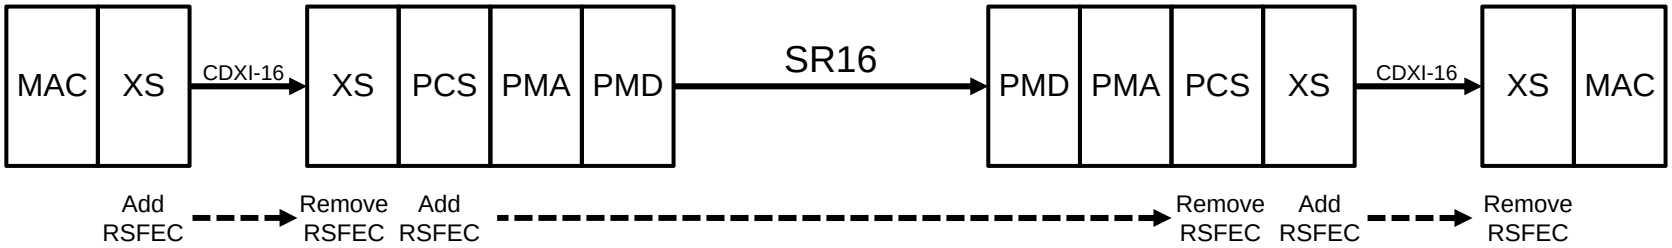

… but this doesn't mean that we are endorsing a 400GBASE-SR16 proposal
Some of the items depicted may be the subject of future projects – some may never exist!

Some notes

- Host systems:
  - Host #1 – multiport MAC ASIC with 16 lane i/f, port extender supports 16 lane module
  - Host #2 – Integrated ASIC driving 16 lane module
  - Host #3 – 1st gen ASIC with 16 lane i/f, 2nd gen port extender supports 8 lane module

- Modules
  - Modules 1x – 16 lane module slots
  - Modules 2x – 8 lane module slots
  - Modules xA – SR16 interface
  - Modules xB – LR4 interface
  - Modules xC – SR8 interface
  - Modules xD – SR4 interface

Config: Host #1, Module 1A

- 1st gen, 16 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

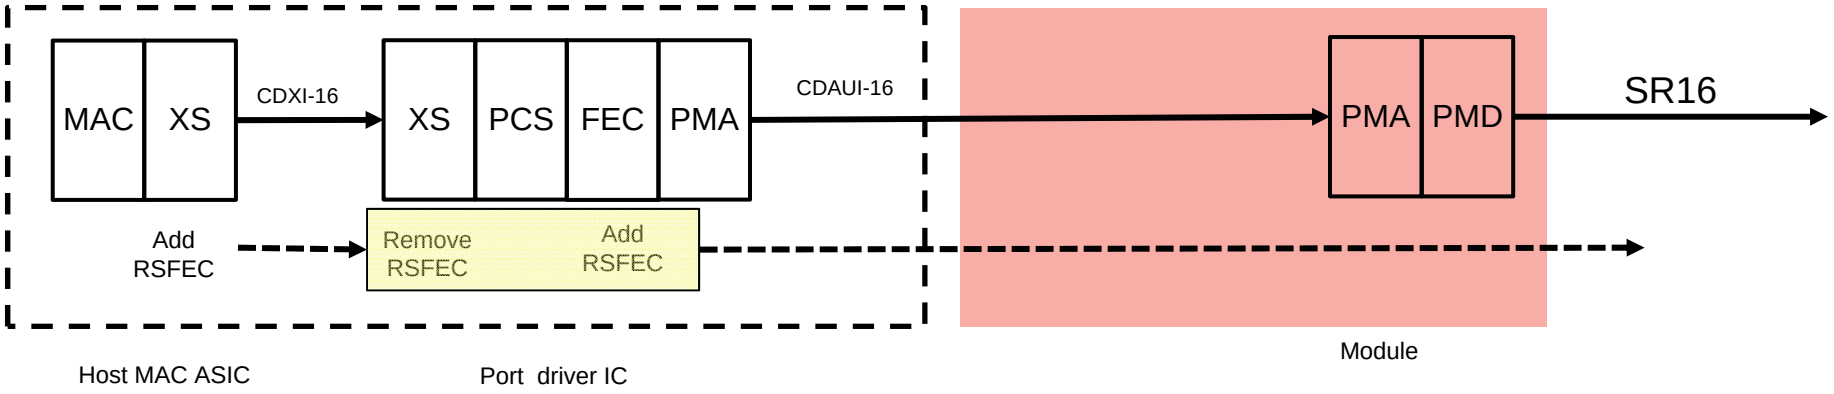

Definition for 16-bit XS & SR16 PCS/FEC/PMA should be as similar as possible so that driver IC has minimal logic

In this configuration, there may be separate FEC statistics for CDXI-16 & CDAUI-16

Config: Host #1, Module 1A (alternate view)

- 1st gen, 16 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

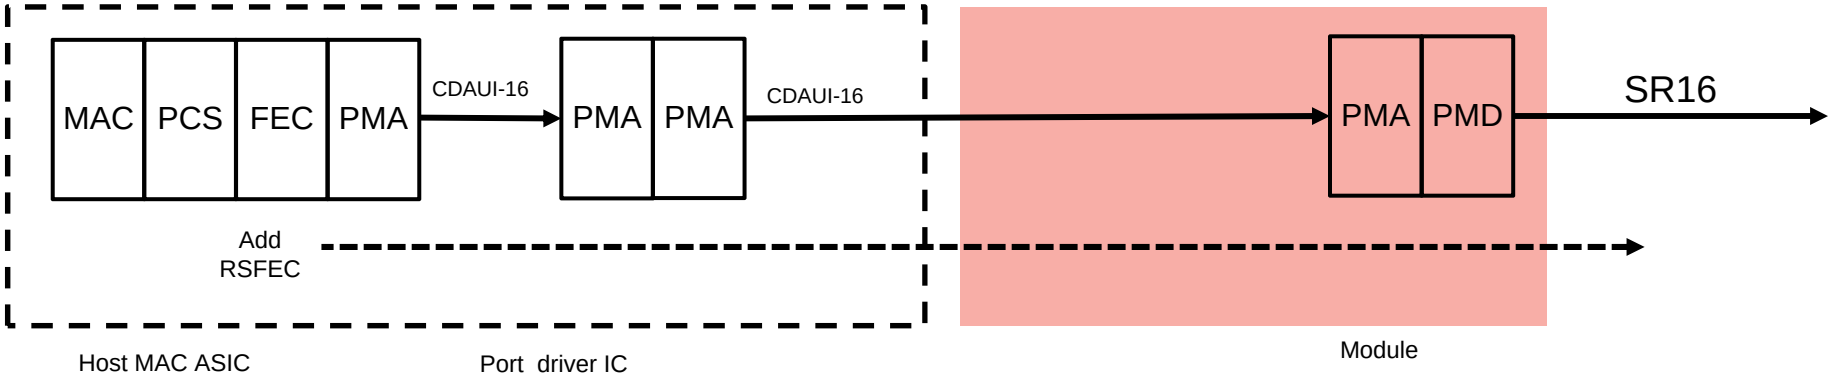

The detailed architecture proposal may require both views, or may make previous view redundant

PMA – PMA function defined for .3ba

Config: Host #1, Module 1B

- 1st gen, 16 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

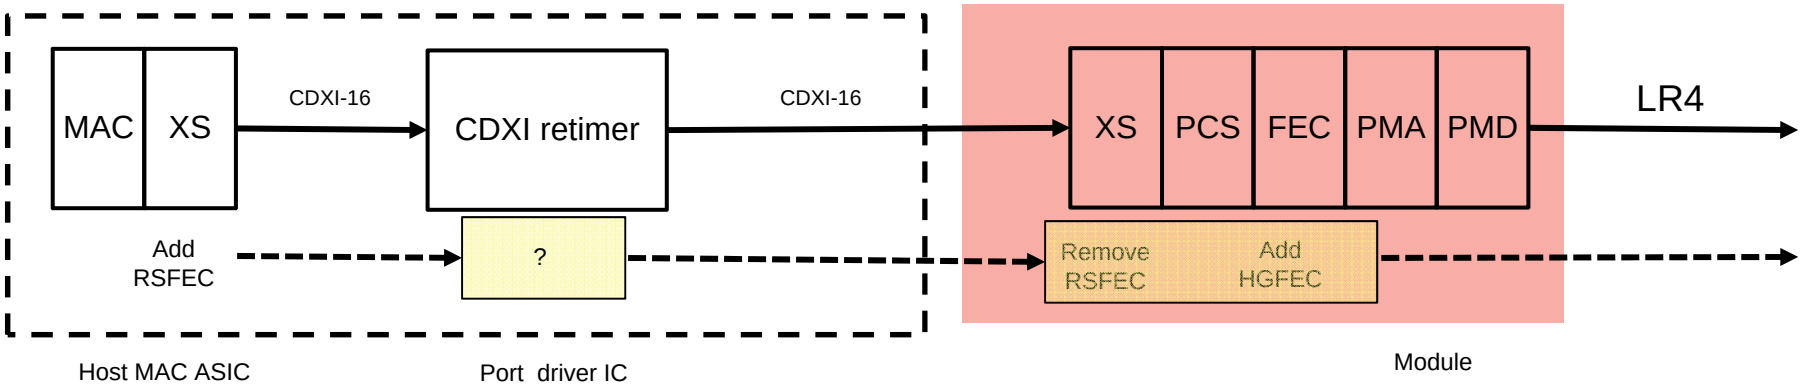

Complex module to handle coding & FEC in advance of host integration

Detailed definition required for CDXI retiming (back-to-back XS with separate FEC statistics; retimer only, etc.)

Config: Host #2, Module 1A

- 2nd gen, 16 lane module interface, integrated host MAC ASIC

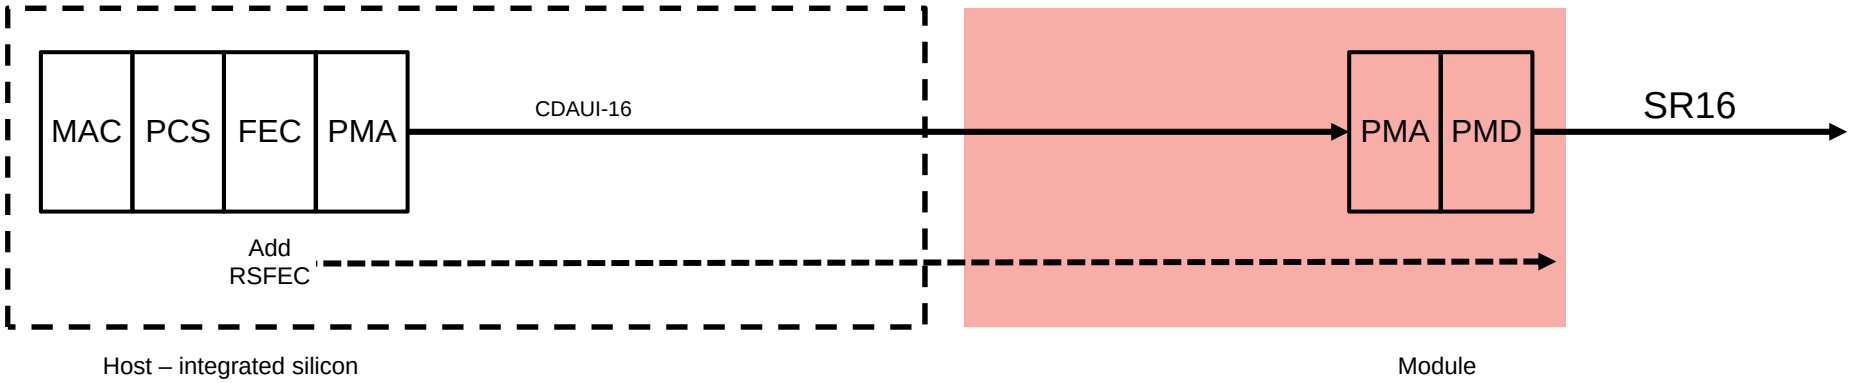

Driver integrated with host ASIC, indistinguishable from config #1 for external view

Config: Host #2, Module 1B

- 2nd gen, 16 lane module interface, integrated host MAC ASIC

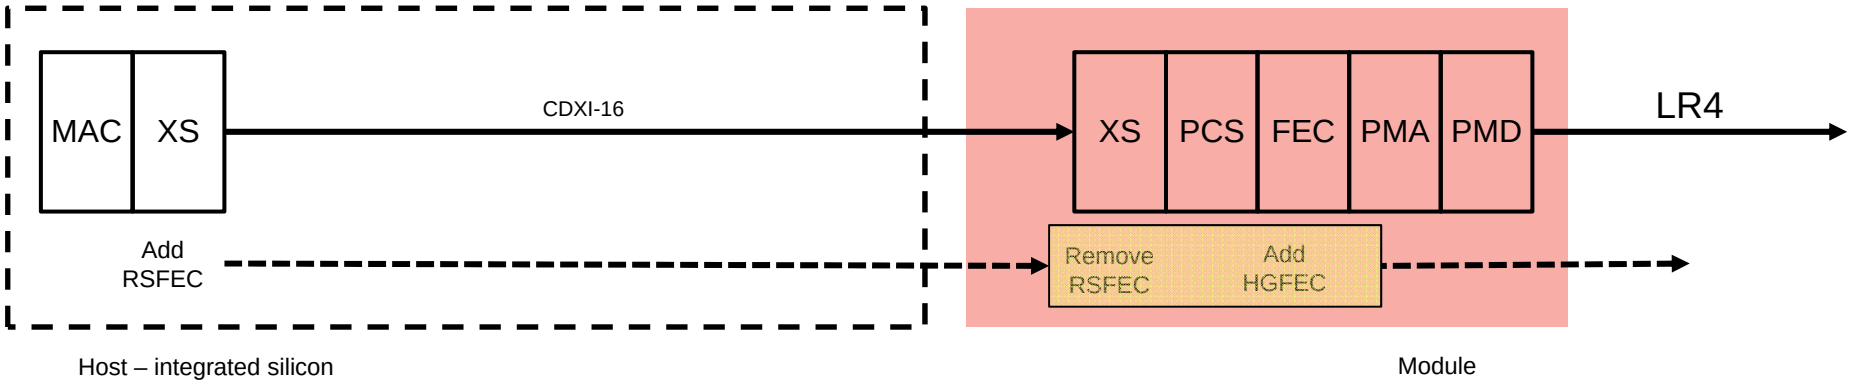

Config: Host #3, Module 2B

- 3rd gen, 8 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

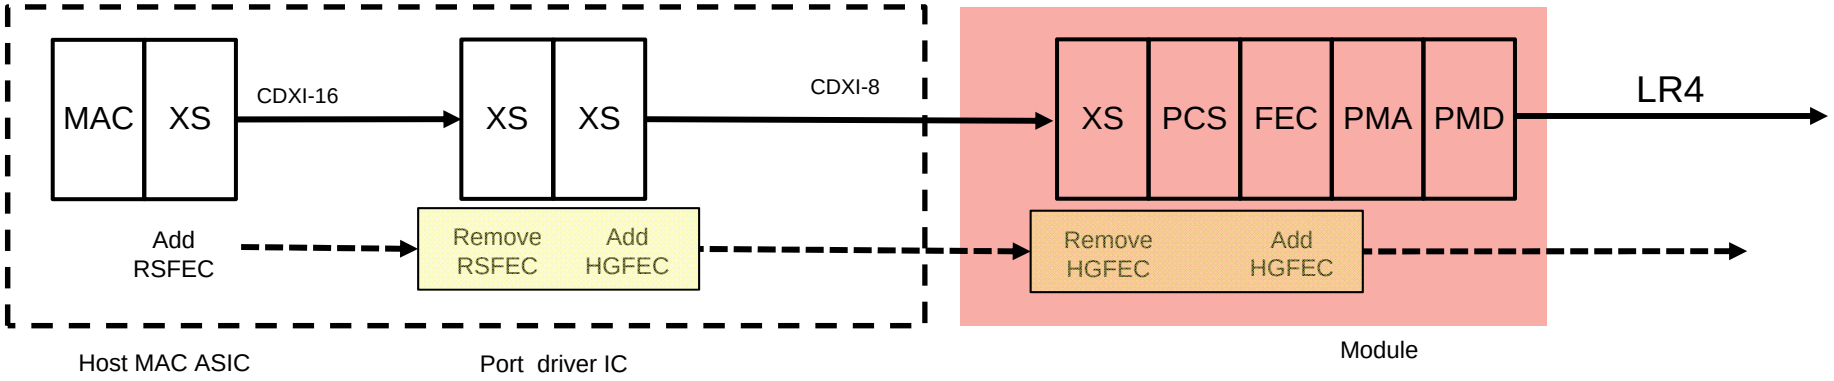

More aggressive C2M interface, re-using 1st gen. host ASIC

If 8x50 C2M requires stronger FEC, back-to-back XS required to change FEC

XS to PCS FEC pass-through to be defined (if 8 lane C2M FEC is similar to 4 lane optical FEC)

Config: Host #3, Module 2B (alternate view)

- 3rd gen, 8 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

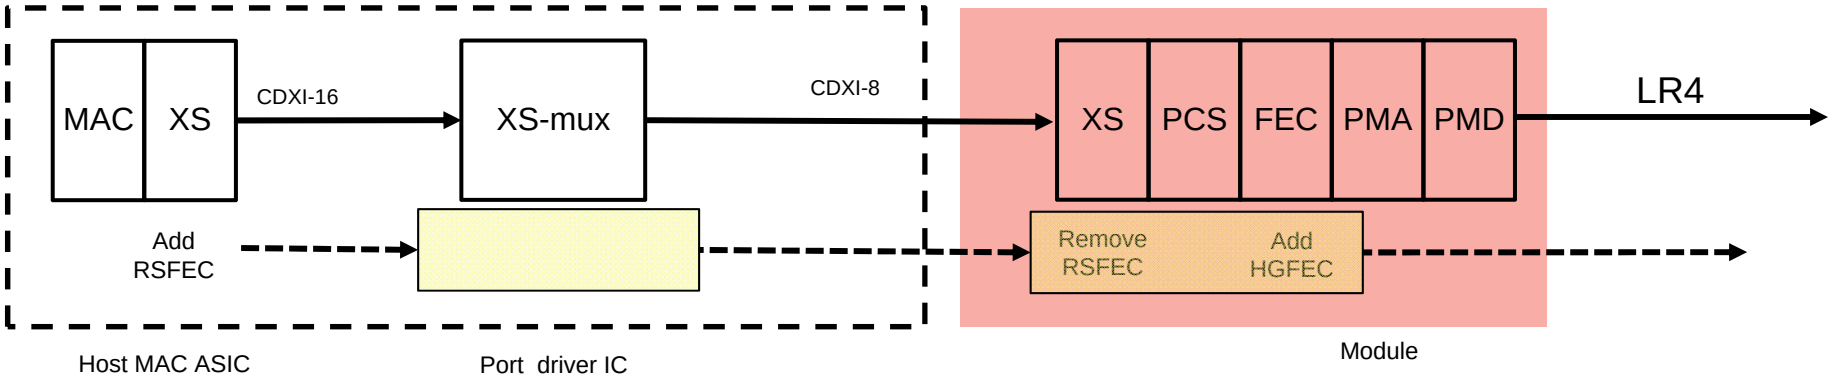

If RS-FEC is sufficient for 8x50 C2M, some form of mux function could change from CDXI-16 to CDXI-8

Config: Host #3, Module 2B (alternate view)

- 3rd gen, 8 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

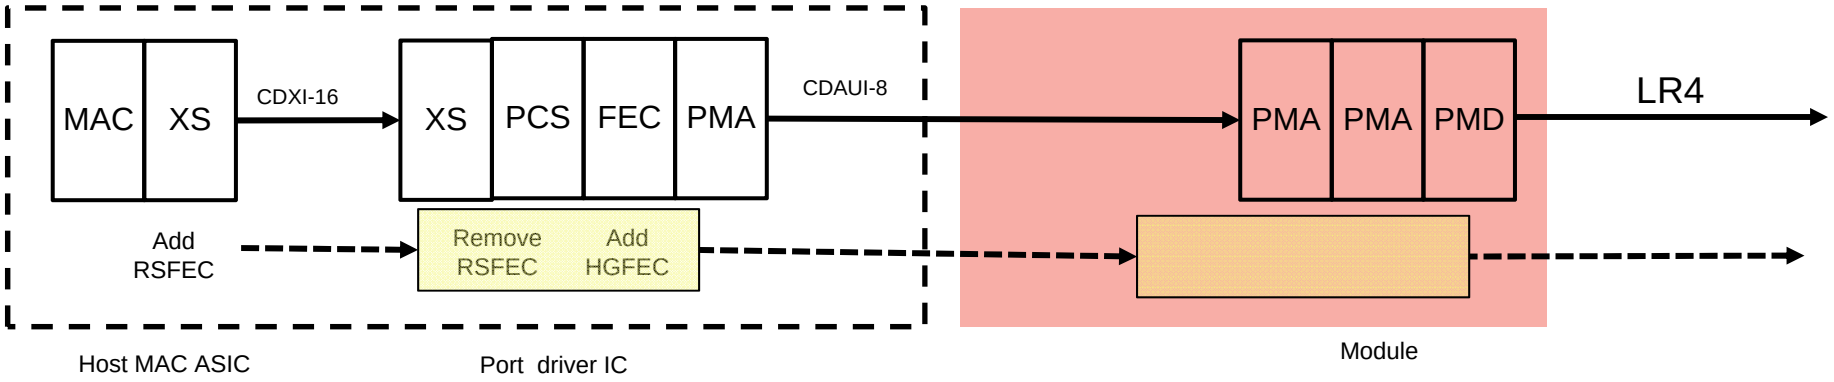

Alternate view could be allowed by the architecture, PMA mux function changes from CDAUI-8 to CDAUI-4

This view requires that the HGFEC used for 4 lane optics is also suitable for 8 lane C2M

Config: Host #3, Module 2A (unclear demand)

- 3rd gen, 8 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

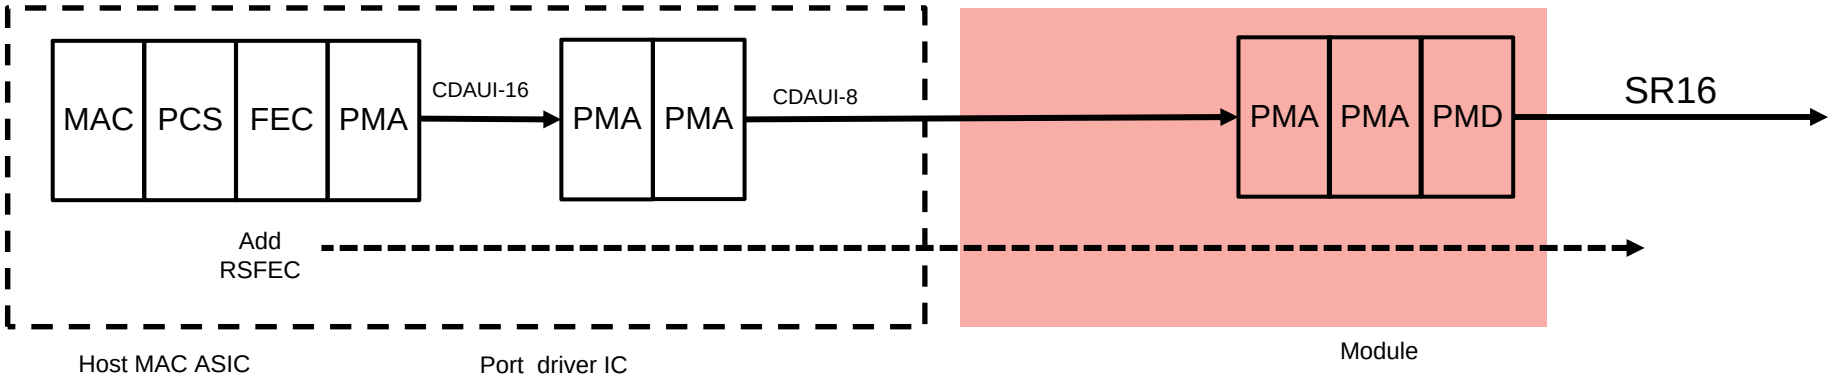

PMA mux from CDAUI-16 to CDAUI-8 & then demux in the module

Generally, interfaces wider than module lane width are uncommon

Config: Host #3, Module 2A (alternate view)

- 3rd gen, 8 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

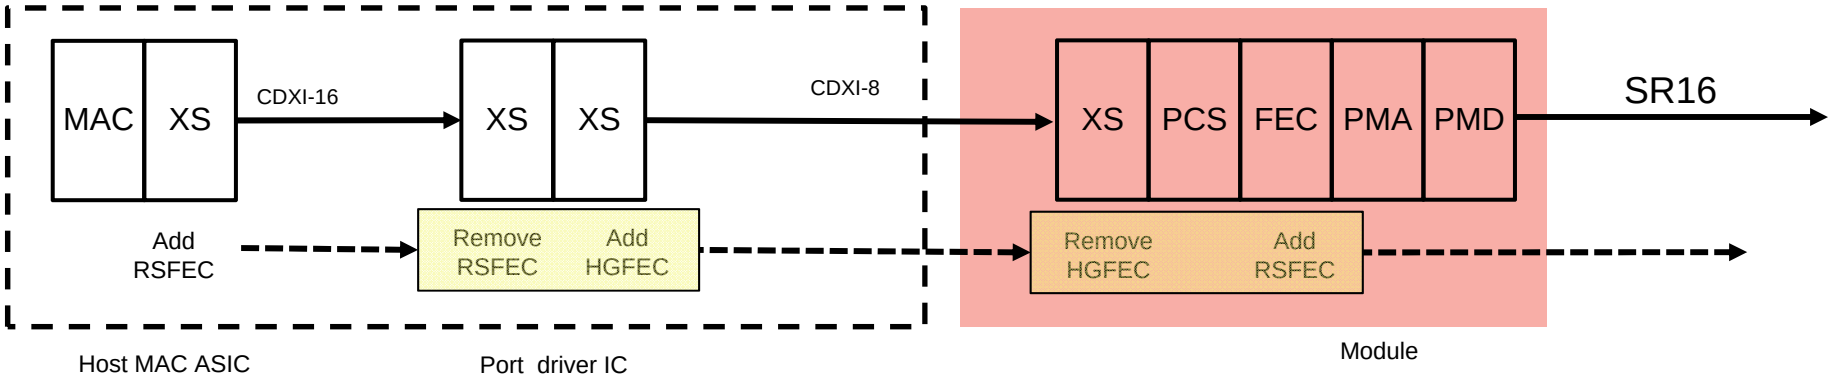

If 8x50 C2M requires stronger FEC, simple SR16 module is not possible

Config: Host #3, Module 2C

- 3rd gen, 8 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

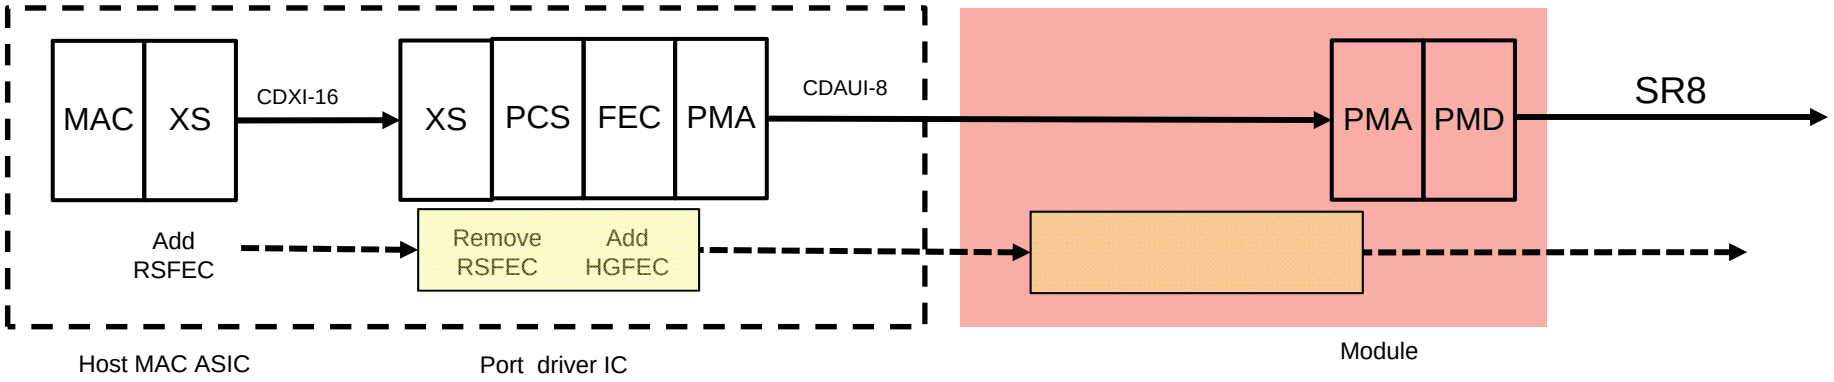

If SR8 uses same HGFEC as 8 lane C2M & 4 lane SMF, then simple 8 lane module is possible

Config: Host #3, Module 2D

- 3rd gen, 8 lane module interface, multi-port host MAC ASIC & extender (port driver) ASIC

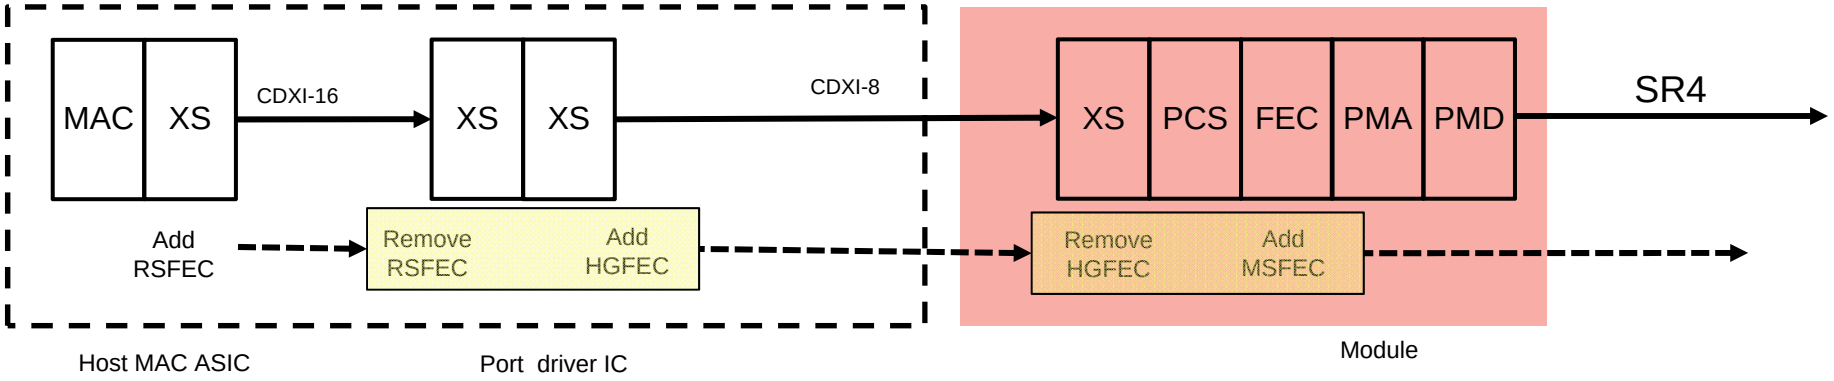

Medium specific FEC for 4 lane MMF – 3 separate FEC segments