100G SR4: FEC & Legacy

Brian Misek, John Petrilla: Avago Technologies
Jonathan King: Finisar
July 2012

Summary

•Various SR and LR4 implementations connecting hosts defined in 802.3ba (BA) and 802.3bj (BJ) are examined for possible conflicts using the following assumptions.

- •CAUI 4NG to be defined in 802.3 Next Gen, incorporates Low Latency FEC (LL FEC) in development in 802.3bj
- •100G LR4 does not require LL FEC for 10 km reach; compatible with CAUI 4
- •100G SR4 requires LL FEC, i.e. CAUI 4NG, for 100 m reach.
- •100G CR4 cases are expected to align with 100G SR4 cases.
- •100G ER4 cases are expected to align with 100G LR4 cases.

•Identified conflicts are due to encoding differences. These can be resolved if BJ hosts and gearboxes can recognize coding conflicts and revert to BA encoding for such connections or other means of auto negotiation are developed .

Terminology

100G LR4 XCVR(CAUI 4): Optical 4-lane, 25G/lane, transparent retiming repeater & transducer
100G LR4 XCVR(CAUI 10): Optical 4-lane, 25G/lane, 10:4 gearbox, retiming repeater & transducer
100G SR4 XCVR(CAUI 4NG): Optical 4-lane, 25G/lane, transparent retiming repeater & transducer
100G SR10 XCVR(CPPI 10): Optical 10-lane, 10G/lane, transparent repeater & transducer
10:4 BA Gearbox: Gears between 802.3ba defined CAUI 4 and CAUI 10
10:4 BJ Gearbox: Terminates BA(BJ) defined link and initiates BJ(BA) link, BJ link encoded for LL FEC
BA Host: Host with PHY defined in 802.3ba
BJ Host: Host with PHY defined in 802.3bj
CAUI 4: 4-lane, 25G/lane, CAUI 4 defined in 802.3ba
CAUI 4NG: 4-lane, 25G/lane ,CAUI 4NG to be defined in 802.3 Next Gen to incorporate 802.3bj defined LL FEC
CAUI 10: 10-lane, 10G/lane, CAUI defined in 802.3ba
CPPI 10: 10-lane, 10G/lane, CPPI defined in 802.3ba
LL FEC: Low Latency FEC defined in 802.3bj
MLG: Multi-Lane Gearbox

100G LR4: Cases without external gearboxes (1)

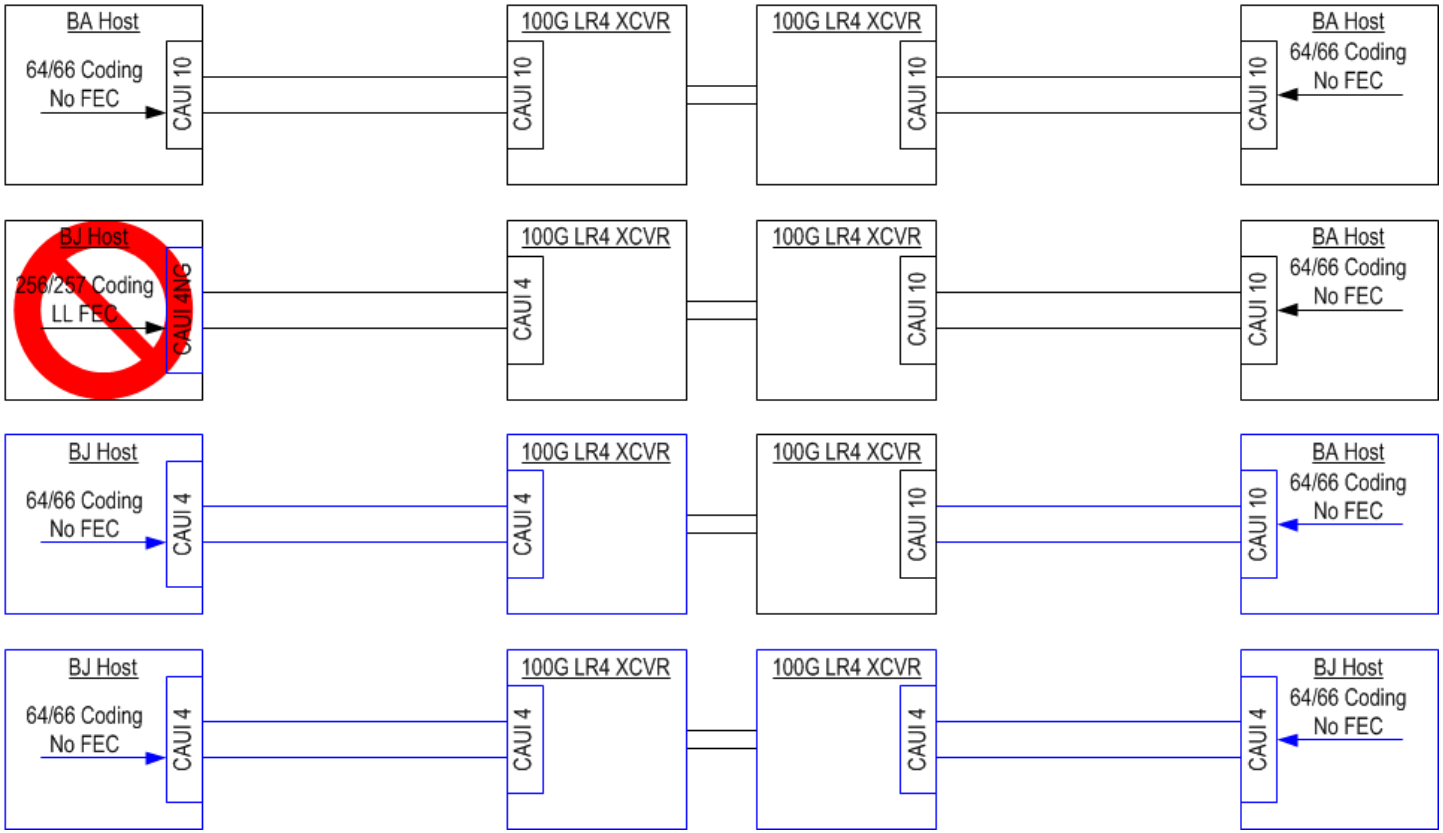

Adopting new generation 100G LR4 and BJ Host IC

- Top block diagram shows existing case where LR4 transceivers with CAUI 10 electrical interfaces connect BA hosts.
- Second block diagram shows an effort to connect a BJ host to a BA host with LR4 transceivers, one with a CAUI 4 electrical interface and one with a CAUI 10. This will not work due to the encoding differences.
- As shown in the third and fourth block diagrams, traffic between a BA host and a BJ host can be supported, if the BJ host recognizes the LR4 connection and reverts to 64/66 encoding with no FEC.

100G LR4: Cases with external gearboxes (2)

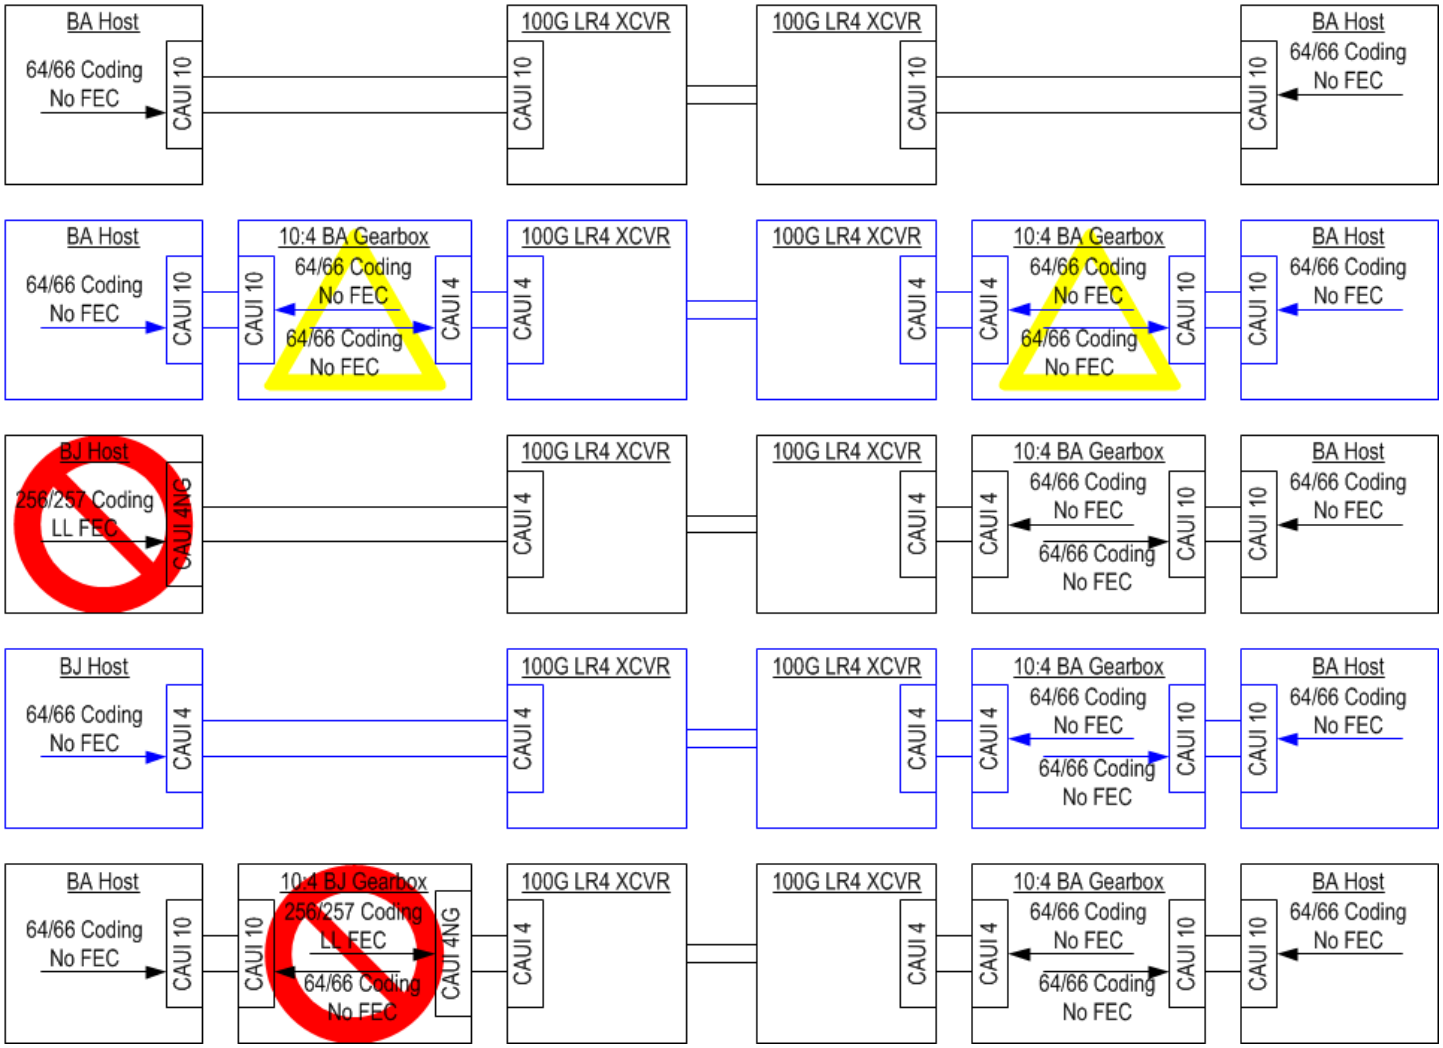

Adopting new generation 100G LR4 with gearboxes

- Top block diagram shows existing case where LR4 transceivers with CAUI 10 electrical interfaces connect BA hosts.
- Second block diagram shows 10:4 BA gearboxes may be used with new generation 100G LR4 transceivers to connect BA hosts. More on caution symbols later.
- Third block diagram shows an effort to connect a BJ host to a BA host with LR4 transceivers with CAUI 4 electrical interfaces that will not work due to the encoding differences.
- Fourth block diagram shows traffic between a BA host and a BJ host can be supported, if the BJ host recognizes the LR connection and reverts to 64/66 encoding with no FEC.
- Bottom block diagram shows a gearbox choice that would inject coding differences unless it reverted to BA coding.

- Traffic between BJ and BA hosts can be supported, if the BJ host recognizes the LR connection and reverts to 64/66 encoding with no FEC. Traffic between BA hosts using new generation LR4 transceivers may be supported with BA or MLG gearboxes if only LR4 transceivers are used in the ports; otherwise, this should be discouraged.

100G SR4: Reuse of 100G SR10 transceivers? (3)

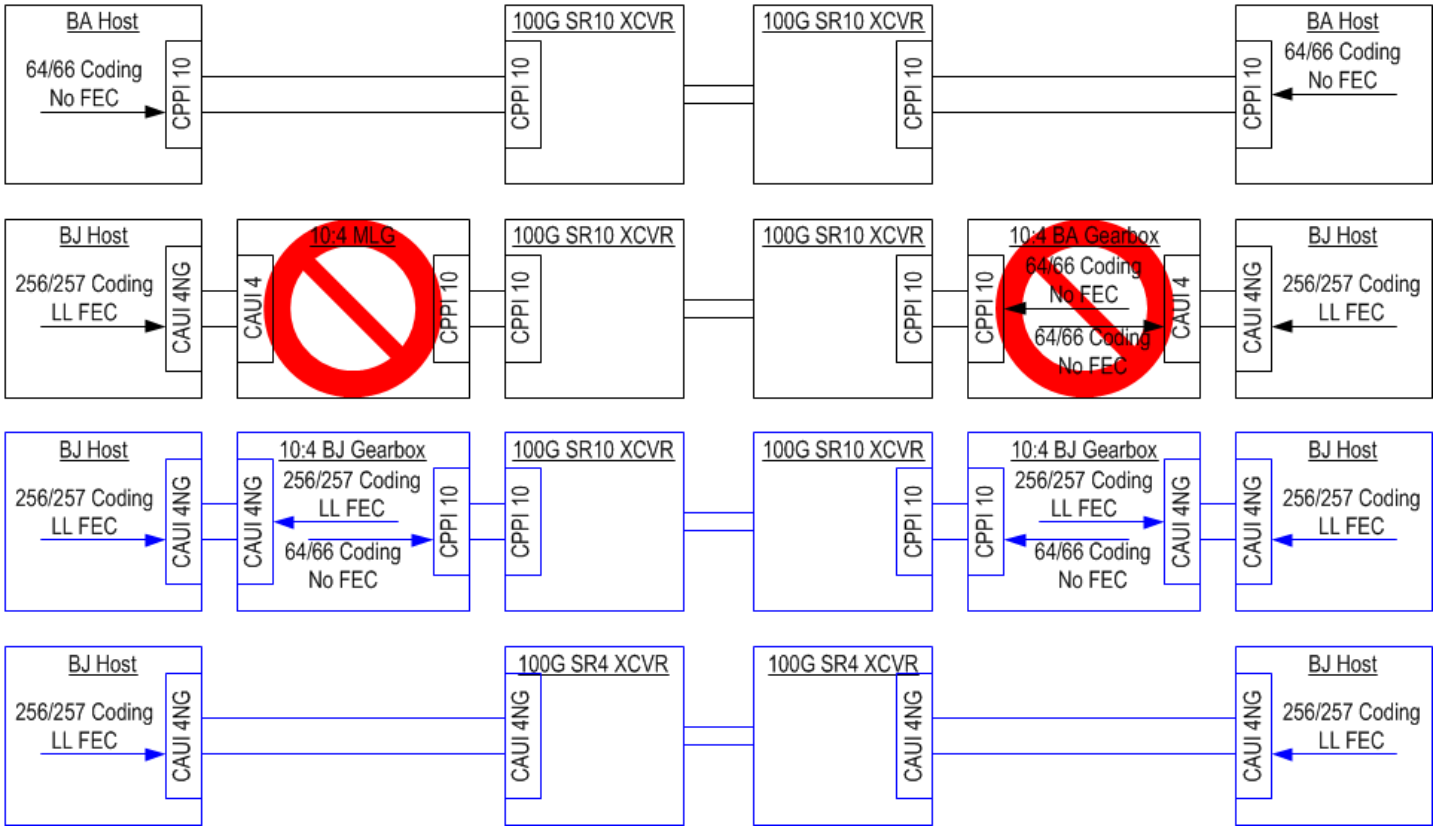

Reuse of 100G SR10 Transceivers with new BJ Host

- Top block diagram shows existing case where SR10 transceivers with CPPI 10 electrical interfaces connect BA hosts.
- Traffic between 4-lane-25G/lane BJ hosts can use BJ specific gearboxes but neither an MLG nor BA gearbox can be used between BA and BJ hosts where LL FEC is active.
- Although gearboxes can be used to connect BJ hosts through SR10 transceivers, such implementations may not be numerous.
- For completeness, the fourth block diagram shows BJ hosts connected via 100G SR4 transceivers as the more likely means of connecting two BJ hosts.

100G SR4: Cases with external gearboxes (4)

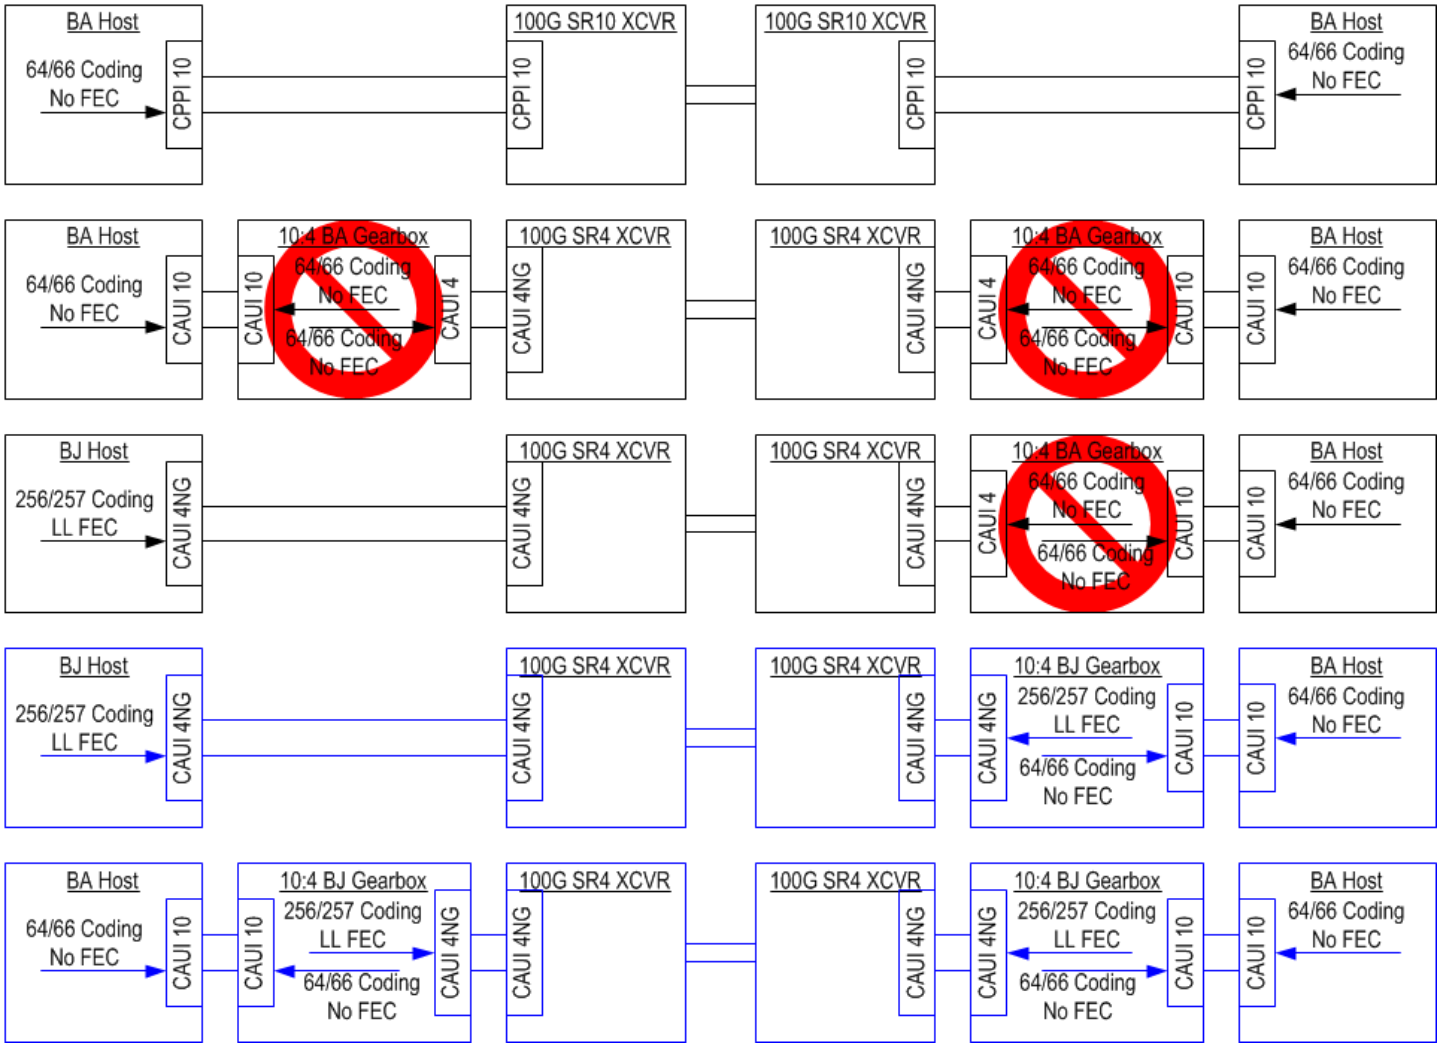

Adopting 100G SR4 and gearboxes

•Top block diagram shows existing case where SR10 transceivers with CPPI 10 electrical interfaces connect BA hosts.

•Second block diagram shows BA gearboxes (could be MLG) used with 100G SR4 transceivers to connect BA hosts that may not support a 100 m MMF or 5 m Cu reach without LL FEC.

• Further, as seen in the third block diagram, this creates a legacy issue for a BJ host combined with 100G SR4.

•A BA host may be connected to a BJ host with a BJ gearbox that terminates the link from one host and recodes the signal for the other while providing LL FEC to support MMF and Cu reach objectives.

•Traffic between 10-lane, 10G/lane BA hosts may be possible with a BA specific gearbox but, since neither MLG nor BA gearboxes can be used between BA and BJ hosts where LL FEC is active or are likely to support 100G CR4, such implementations should be discouraged.

100G: Common LR4 - SR4 ports (5)

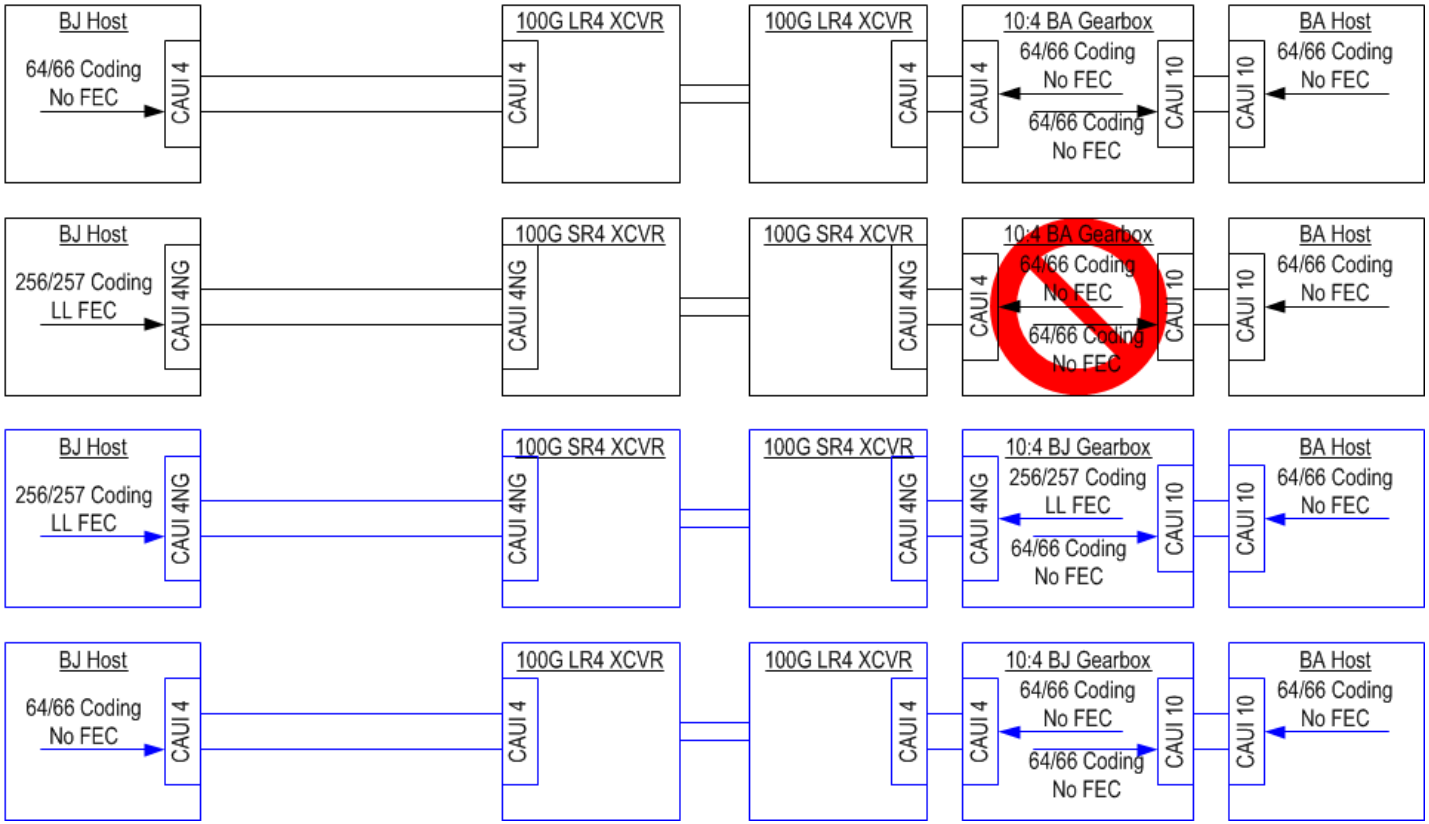

Common ports for 100G SR4 and LR4 Transceivers with BA and BJ hosts

- Top block diagram shows case where LR10 transceivers connect a BJ host to a BA host with a BA gearbox. As shown in the second block diagram such a BA gearbox is not sufficient for a port that accepts both SR4 and LR4 transceivers.
- As shown in the third and fourth block diagrams, traffic between a BA host and a BJ host through ports that can support both LR4 and SR4 can be supported, if the BJ gearbox and host recognize the LR connections and revert to 64/66 encoding with no FEC.

- Conflicts between BA and BJ hosts and gearboxes can be resolved, if BJ hosts and gearboxes can recognize LR4 connections and revert to BA mode when appropriate.

Should 802.3bj LL FEC be used with 100G SR4 Transceivers?

- 802.3bj appears likely to incorporate Low Latency FEC (LL FEC)
- Equipment designers, e.g. early adopters, may want to take advantage of the port density offered by 100G SR4 and new LR4 form factors before host ICs implementing 802.3bj (BJ) are available, e.g. using a 10-lane-10G/lane 802.3 ba (BA) host IC with a 10:4 gearbox.
- It is also expected that there is interest in aggregation of 10G as well as 10G/lane optical traffic into 25G/lane traffic.
- Gearboxes , e.g. a MLG or a BA 10:4 PMA, are expected to be available and sufficient to support gearing between 10G/lane and 25G/lane traffic but supporting10G traffic to/from 25G/lane traffic seems problematic.
- Since, BJ 25G/lane traffic with Low Latency FEC (LL FEC) (see gustlin_01_0312_NG100GOPTX) activated does not permit simple lane count changes, BA level gearboxes are no longer sufficient for BJ traffic and a BA:BJ (BJ) gearbox will be needed.
- Questions to be considered include:
    - If optical channel reach objectives can be supported without using LL FEC, should LL FEC be activated for optical channels? There may be ports supporting both optics and copper cables where LL FEC is needed for the copper cable reach – can this be accommodated?
    - Should optical channel reach be defined for both cases; where LL FEC is activated and where it is not? Longer reaches are expected for FEC encoded traffic – how would users deal with different reaches using the same module that depends on FEC activation?
    - If optical channel reach objectives require LL FEC or there are significant cost and power advantages for LL FEC traffic:
      - should early adopters depend on availability of BJ gearboxes?
      - should the upcoming Next Gen task force not take advantage of LL FEC and give priority to simple gearboxes over cost and power savings of optical transceivers?
    - Is this an issue for BJ, the Next Gen study group and/or upcoming Next Gen task force to consider or best left to implementers?
- Are there other issues to consider?

BACKUP

100GE: BA, BJ, Next Gen legacy issues without gearboxes

	Optical Interface	Form Factor	Electrical Interface	Legacy Issue
802.3ba	100GBASE-SR10	CXP	CPPI 10	None – 10-lane optical links only expected for BA, not carried to Next Gen
802.3ba	100GBASE-SR10	CFP/CFP2	CAUI 10	None – 10-lane optical links only expected for BA, not carried to Next Gen
802.3ba	100GBASE-LR4	CFP	CAUI 10	Conflict if connected to LR4 with CAUI 4NG interface using 256/257 & LL FEC
802.3ba	100GBASE-ER4	CFP	CAUI 10	Conflict if connected to ER4 with CAUI 4NG interface using 256/257 & LL FEC
802.3bj	NA	QSFP	100GBASE-CR4	None – No legacy 100G Copper links in BA
802.3NG	100GBASE-SR4	QSFP	CAUI 4NG	None – No legacy 100G 4-lane SR links in BA
802.3NG	100GBASE-SR4	CFP2/4	CAUI 4NG	None – No legacy 100G 4-lane SR links in BA
802.3NG	100GBASE-?R?	QSFP	CAUI 4NG	(?R? for 500 m SMF) None – No legacy 100G ?R? links in BA
802.3NG	100GBASE-?R?	CFP2/4	CAUI 4NG	None – No legacy 100G ?R? links in BA
802.3NG	100GBASE-LR4	CFP2/4	CAUI 4NG	Conflict if connected to LR4 with CAUI 10 interface
802.3NG	100GBASE-ER4	CFP2/4	CAUI 4NG	Conflict if connected to ER4 with CAUI 10 interface

100GE: BA, BJ, Next Gen legacy issues with gearboxes

	Optical Interface	Form Factor	Electrical Interface	Legacy Issue
802.3ba	100GBASE-SR10	CXP	CPPI 10	Conflict if CPPI 10 geared to 4 lanes for Host BJ interface without terminating CAUI link, coding to 256/257 and adding LL FEC
802.3ba	100GBASE-SR10	CFP	CAUI 10	Conflict if CPPI 10 geared to 4 lanes for Host BJ interface without terminating CAUI link, coding to 256/257 and adding LL FEC
802.3ba	100GBASE-LR4	CFP	CAUI 10	Conflict if CAUI 10 geared to 4 lanes encounters CAUI 4NG with 256/257 & active LL FEC
802.3ba	100GBASE-ER4	CFP	CAUI 10	Conflict if CAUI 10 geared to 4 lanes encounters CAUI 4NG with 256/257 & active LL FEC
802.3bj	NA	QSFP	100GBASE-CR4	None – No legacy 100G Copper links in BA
802.3NG	100GBASE-SR4	QSFP	CAUI 4	Conflict if CAUI 4NG is geared to 10 lanes for BA interface without terminating CAUI 4NG link, removing LL FEC and coding to 64/66
802.3NG	100GBASE-SR4	CFP2/4	CAUI 4	Conflict if CAUI 4NG is geared to 10 lanes for BA interface without terminating CAUI 4NG link, coding to 64/66 and removing LL FEC
802.3NG	100GBASE-?R?	QSFP	CAUI 4	(?R? for 500 m SMF) None – No legacy 100G ?R? SMF links in BA
802.3NG	100GBASE-?R?	CFP2/4	CAUI 4	None – No legacy 100G ?R? links in BA
802.3NG	100GBASE-LR4	CFP2/4	CAUI 4	Conflict if CAUI 4NG LR4 with 256/257 & active LL FEC connects to LR4 with CAUI 10 electrical interface or CAUI 10 geared to 4 lane interface
802.3NG	100GBASE-ER4	CFP2/4	CAUI 4	Conflict if CAUI 4 LR4 with 256/257 & active LL FEC connects to LR4 with CAUI 10 electrical interface or CAUI 10 geared to 4 lane interface