

BER Objectives for Physical Layers using 400 Gb/s per lane interfaces

Matt Brown, Qualcomm

David Ofelt, HPE

Kent Lusted, Synopsys

Supporters

- Adam Healey, Broadcom
- Jeff Slavick, Broadcom
- Xiang He, Huawei
- Guangcan Mi, Huawei
- Mike Dudek, Marvell
- Shawn Nicholl, AMD
- Leesa Noujeim, Google
- John D'Ambrosia, FutureWei

Introduction

- An objective for the error ratios at the MAC service interface is useful to provide guidance on the performance targets for 400 Gb/s/lane (400GPL) interfaces.
- This contribution provides some background including constraints set by objectives and error definitions and allocations defined in previous projects.
- Finally, related objectives for the 400GPL project are proposed.

BER objectives in previous projects

802.3bs (200 Gb/s and 400 Gb/s Ethernet) BER objective slides

https://www.ieee802.org/3/bs/Objectives_16_0317.pdf

Project Objectives

- Support a MAC data rate of 200 Gb/s
- Support a MAC data rate of 400 Gb/s
- Support a BER of better than or equal to 10^{-13} at the MAC/PLS service interface (or the frame loss ratio equivalent)
- Support full-duplex operation only
- Preserve the Ethernet frame format utilizing the Ethernet MAC
- Preserve minimum and maximum FrameSize of current Ethernet standard

802.3df (800 Gb/s Ethernet) BER objective slides

https://www.ieee802.org/3/df/proj_doc/objectives_P802d3df_221117.pdf

Adopted IEEE P802.3df Objectives

Non-Rate Specific

- Support full-duplex operation only
- Preserve the Ethernet frame format utilizing the Ethernet MAC
- Preserve minimum and maximum FrameSize of current IEEE 802.3 standard
- Support a BER of better than or equal to 10^{-13} at the MAC/PLS service interface (or the frame loss ratio equivalent)
- Provide support to enable mapping over OTN

802.3dj (1.6 Tb/s Ethernet BER objective slides)

https://www.ieee802.org/3/dj/projdoc/objectives_P802d3dj_240314.pdf

Adopted IEEE P802.3dj Objectives (1 of 2)

Non-Rate Specific

- Support full-duplex operation only
- Preserve the Ethernet frame format utilizing the Ethernet MAC
- Preserve minimum and maximum FrameSize of current IEEE 802.3 standard
- Support a BER of better than or equal to 10^{-13} at the MAC/PLS service interface (or the frame loss ratio equivalent)
- Provide support to enable mapping over OTN

Effective BER target is $1E^{-13}$ at the MAC service interface for 400GE, 800GE, and 1.6TE

There is a possibility that a higher Ethernet (e.g., 3.2TE or 6.4TE) will have a different BER target but that is beyond our ability to predict at this time.

These previous objectives inherently already impose a constraint on new Physical Layer implementations using 400GPL.

Frame loss ratio

- Use of FEC in the PCS means that after PCS decoding, including error correction, either there are no singular errors or there is a large set of bits that are discarded when an FEC codeword is uncorrectable.
- As such, the effective BER target of $1\text{E-}13$ has been translated into a frame loss ratio (FLR) target at the MAC service interface as follows:
 - “The frame loss ratio for 64-octet MAC frames with minimum interpacket gap is expected to be less than 6.2×10^{-11} .”

Formalization of allocations and error correlation constraints

- IEEE 802.3dj (currently Draft 3.0) formalizes the allocation of error ratio to a MAC-to-MAC path. See Annex 174A.
 - It allocates a small portion for the FLR ($0.1E-11$) to each of the xMII Extenders leaving $6.0E-11$ for the PCS-to-PCS path.
 - BER and FLR is allocated to the various physical interfaces according to the table and figures that follow.
 - These allocations can be readjusted inside of a system where the equipment is controlled/engineered. For instance, a PHY using co-packaged optics (CPO) could allocate 100% BER to the PMD.
 - Provides methodology and metric to constrain the correlation of errors with the parameter block error ratio (BLER) to ensure the target FLR is achieved.

IEEE 802.3dj Annex 174A FLR allocations

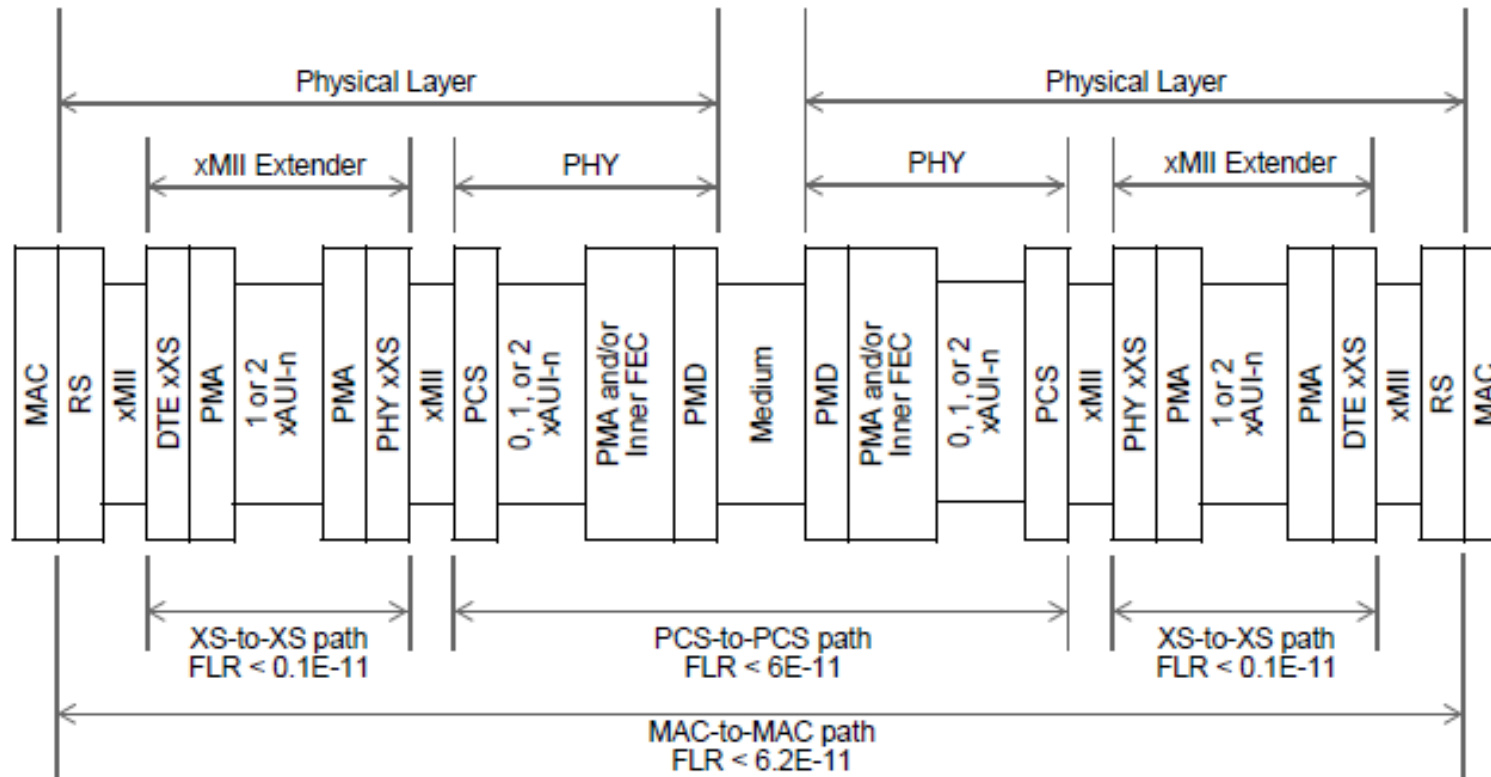


Figure 174A-6—Frame loss ratio allocations for a MAC-to-MAC path

IEEE 802.3dj Annex 174A FLR/BER allocations

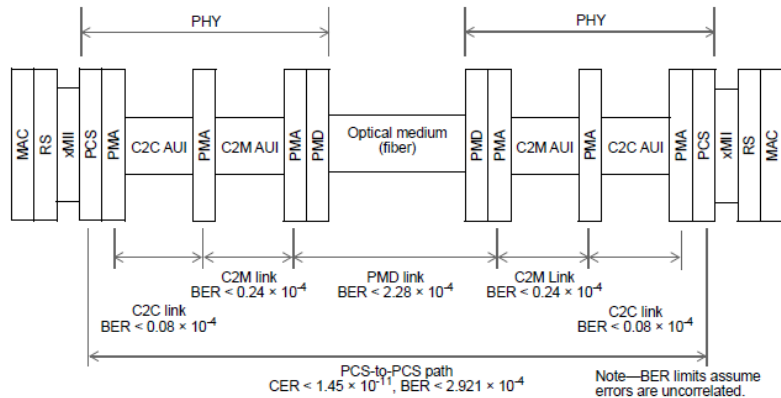


Figure 174A-9—Error allocations for optical PHY types with no FEC sublayer

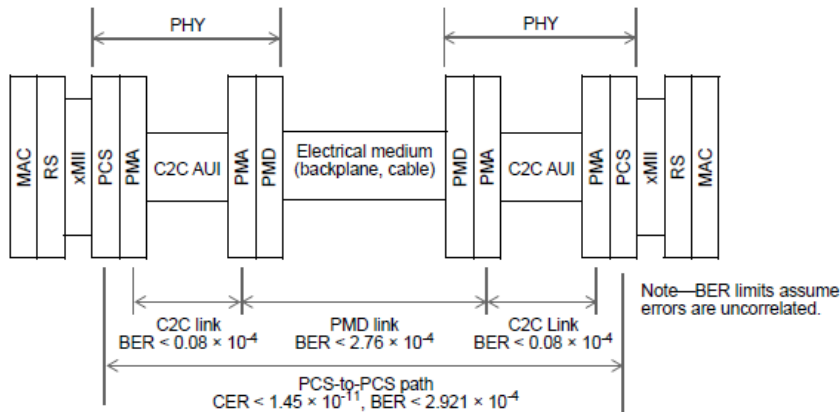


Figure 174A-8—Error allocations for electrical PHY types

Table 174A-1—Error ratio allocations for optical PHYs with no FEC sublayer or with an Inner FEC sublayer

ISL	Frame loss ratio for entire PCS-to-PCS path	Codeword error ratio for entire PCS-to-PCS path	BER for entire PCS-to-PCS path (BER_{total})	BER per ISL ^a
xAUI-n C2C ^b	6×10^{-11}	1.45×10^{-11}	2.921×10^{-4}	0.08×10^{-4}
xAUI-n C2M				0.24×10^{-4}
PMD-to-PMD				2.28×10^{-4}
xAUI-n C2M				0.24×10^{-4}
xAUI-n C2C ^b				0.08×10^{-4}

^a Measured at the PMA closest to the AUI component or PMD and after Inner FEC decoding, if present, except measured at the Inner FEC for 800GBASE-LR1.

^b If the PMD is a type defined in Clause 180, Clause 181, Clause 182, Clause 183, or Clause 185, and xAUI-n C2C is a type defined in Annex 120D (i.e., 50 Gb/s per lane) or Annex 120F (i.e., 100 Gb/s per lane), the xAUI-n C2C is expected to meet the BER allocations in this table.

Table 174A-2—Error ratio allocations for electrical PHYs

ISL	Frame loss ratio for entire PCS-to-PCS path	Codeword error ratio for entire PCS-to-PCS path	BER for entire PCS-to-PCS path (BER_{total})	BER per ISL ^a
xAUI-n C2C ^b	6×10^{-11}	1.45×10^{-11}	2.921×10^{-4}	0.08×10^{-4}
PMD-to-PMD				2.76×10^{-4}
xAUI-n C2C ^b				0.08×10^{-4}

^a Measured at the PMA closest to the AUI component or PMD.

^b If the PMD is a type defined in Clause 178 or Clause 179 (i.e., 200 Gb/s per lane) and xAUI-n C2C is a type defined in Annex 120D (i.e., 50 Gb/s per lane) or Annex 120F (i.e., 100 Gb/s per lane), the xAUI-n C2C is expected to meet the BER allocations in this table.

Legacy considerations

- We need to be considerate of interoperation within networks mixing 400GPL and previous generations of interfaces.
- Consider the following two cases:
 - A 400GPL PMD should work with legacy, e.g., 200GPL C2M and/or C2C AUI
 - the error allocation should be no worse than a 200GPL PMD
 - A 400GPL AUI should work with legacy, e.g., 200GPL PMD and AUIs
 - the error allocation should be no worse than a 200GPL AUI

Summary

- The BER and FLR expectations at the MAC service interface should be the same as already defined for 400GE, 800GE, and 1.6TE.
 - $BER < 1E^{-13}$, $FLR < 6.2E^{-11}$
- To be considerate of legacy interoperability...
 - The PMD error allocation should be consistent with error allocation for a similar PMD as defined in 802.3dj.
 - The C2M AUI error allocation should be consistent with error allocation for a C2M AUI defined in 802.3dj.
 - The C2C AUI error allocation should be consistent with error allocation for a C2C AUI defined in 802.3dj.
 - However, we need not constrain these explicitly with objectives but rather should keep these in mind when developing baselines specifications later.
- There may be other applications or other network configurations (other paradigms) that we have not yet considered, but we can address those later.

Proposed objectives

- We propose the following related objective:
 - Support a BER of better than or equal to 10^{-13} (or the frame loss ratio equivalent) at the MAC/PLS service interface.

Thanks!