Delay, A History of Defining Maximum Latency

(Optical Module Delay)

Jeffery Maki
Juniper Networks

MAC Control PAUSE – Need for Delay Spec

IEEE Std 802.3-2022, IEEE Standard for Ethernet
SECTION THREE

36.5 Delay constraints

In half duplex mode, proper operation of a CSMA/CD LAN demands that there be an upper bound on the propagation delays through the network. This implies that MAC, PHY, and repeater implementations conform to certain delay minima and maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices. MAC constraints are contained in 35.2.4 and Table 35–5. Topological constraints are contained in Clause 42.

In full duplex mode, predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) also demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementations conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices.

The reference point for all MDI measurements is the 50% point of the mid-cell transition corresponding to the reference bit, as measured at the MDI.

Delay Sublayers

PMA Stages

3 PMA Stages

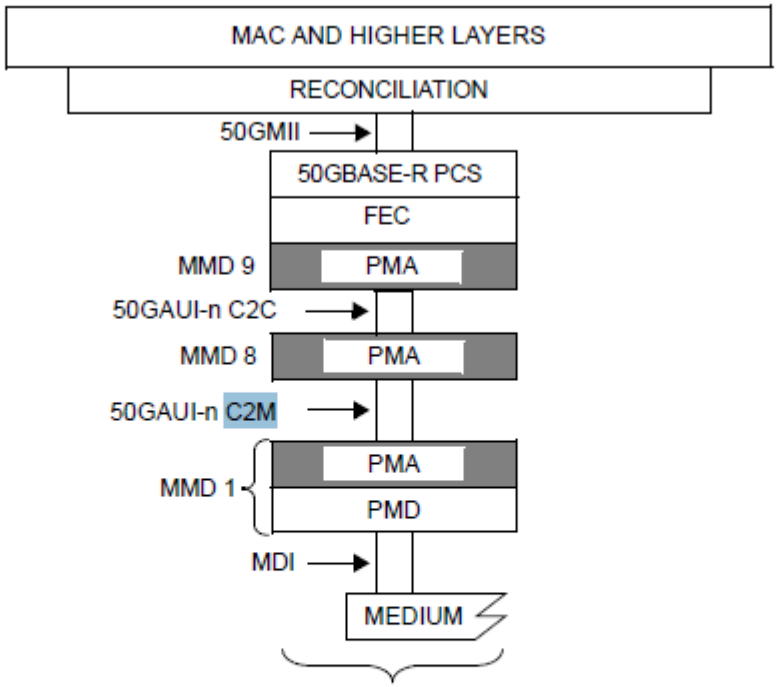

50GAUI-n = 50 Gb/s ATTACHMENT UNIT INTERFACE
50GMII = 50 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

Figure 135A–4—Intermediate PMA device for module interface, FEC implemented with PCS

4 PMA Stages

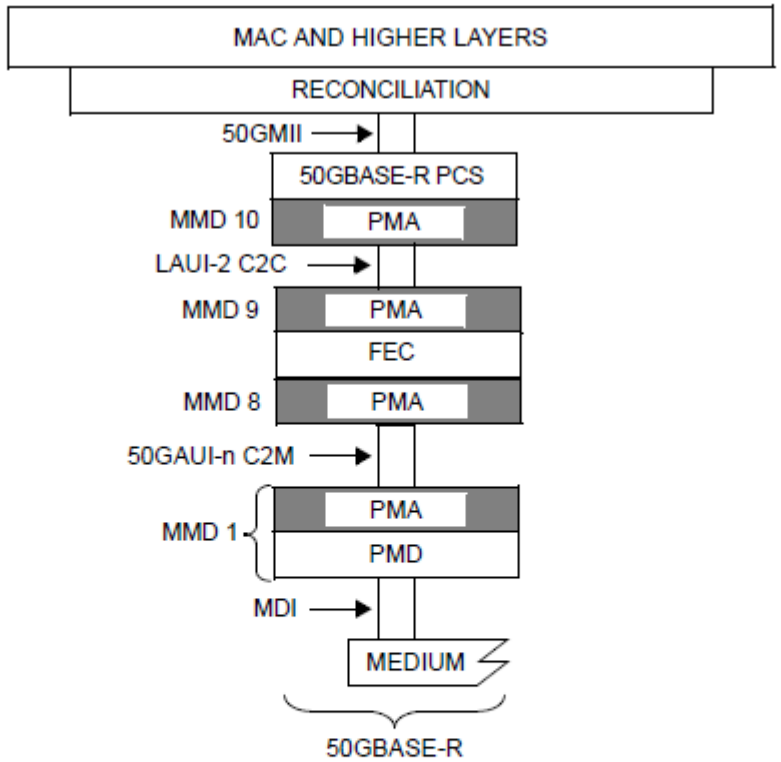

50GAUI-n = 50 Gb/s ATTACHMENT UNIT INTERFACE
50GMII = 50 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
LAUI-2 = 50 Gb/s ATTACHMENT UNIT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

Figure 135A–5—Intermediate PMA device with FEC for module interface

4 PMA Stages

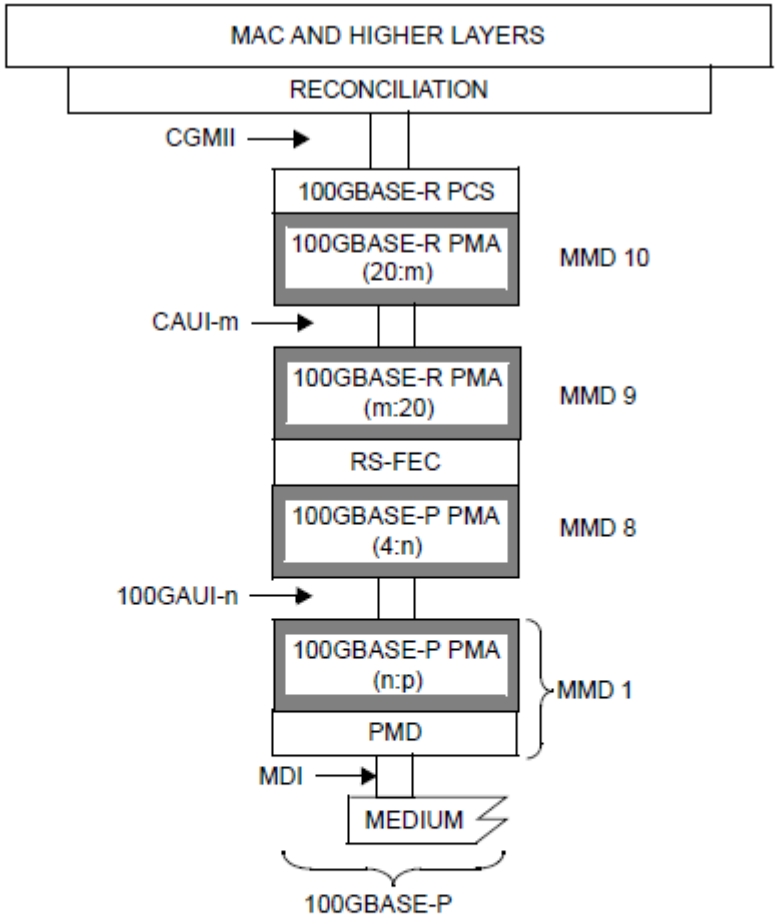

CAUI-10 = 100 Gb/s TEN-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MMD = MDIO MANAGEABLE DEVICE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION
m = 4 or 10
n = 2 or 4
p = 1 or 2

Figure 135A–6—Example 100GBASE-P PHY with CAUI-n and 100GAUI-n

5 PMA Stages

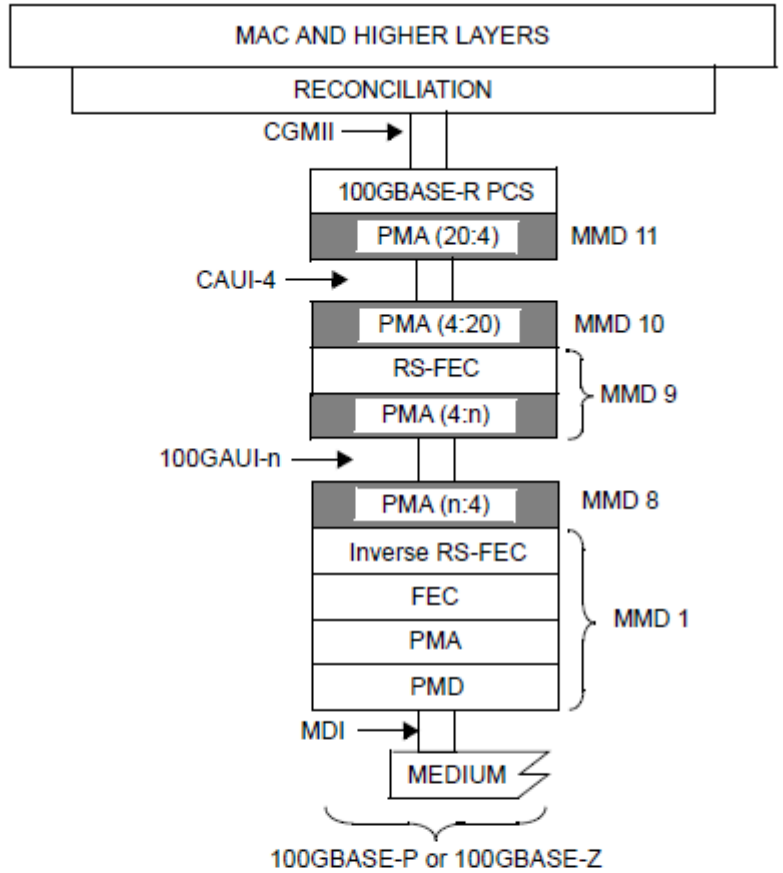

100GAUI-n = 100 Gb/s n-LANE ATTACHMENT UNIT INTERFACE
CAUI-4 = 100 Gb/s FOUR-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MMD = MDIO MANAGEABLE DEVICE
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

Figure 135A–8—Example CAUI-4 chip-to-chip and 100GAUI-n chip-to-module with Inverse RS-FEC

40G, 50G, and 100G – Is it 3 or 4 Stages?

83.5.4 Delay constraints

The maximum cumulative delay contributed by up to four PMA stages in a PHY (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 83–1. A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 80.4 and its references.

Table 83–1—Delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
40GBASE-R PMA	4096	8	102.4
100GBASE-R PMA	9216	18	92.16

What is a Stage?

135.5.4 Delay constraints

The maximum cumulative delay contributed by up to three PMA stages in a PHY (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 135–1. A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 131.4 and its references.

Table 135–1—Delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
50GBASE-R PMA	4608	9	92.16
100GBASE-P PMA	9216	18	92.16

25G, 200G, 400G and 800G – Is 4 Stages Correct?

109.5 Delay constraints

The maximum cumulative delay contributed by up to four PMA stages in a PHY (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 109–1. A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 105.5.

Table 109–1—Delay constraints

Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
4096	8	163.84

120.5.4 Delay constraints

The maximum cumulative delay contributed by up to four PMA stages in a PHY (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 120–1. A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 116.4 and its references.

Table 120–1—Delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
200GBASE-R PMA	18 432	36	92.16
400GBASE-R PMA	36 864	72	92.16

173.4.4 Delay constraints

The maximum cumulative delay contributed by up to four PMA stages in a PHY (sum of transmit and receive delays at one end of the link) shall meet the values specified in Figure 173–1. A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 169.4 and its references.

Table 173–1—Delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
800GBASE-R PMA	73 728	144	92.16

40G

- 40G Module Delay
  - PMA: 102.4 ns / 4
  - Optical PMD: 25.6 ns
  - Total: 51.2 ns

Table 80–7—Sublayer delay constraints

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
40G MAC, RS, and MAC Control	16 384	32	409.6	See 81.1.4.
40GBASE-R PCS	11 264	22	281.6	See 82.5.
40GBASE-R FEC	24 576	48	614.4	See 74.6.
40GBASE-R PMA	4 096	8	102.4	See 83.5.4.
40GBASE-T PHY	25 600	50	640	See 113.11
40GBASE-KR4 PMD	2 048	4	51.2	Includes delay of one direction through backplane medium. See 84.4.
40GBASE-CR4 PMD	4 096	8	102.4	Does not include delay of cable medium. See 85.4.
40GBASE-SR4 PMD	1 024	2	25.6	Includes 2 m of fiber. See 86.3.1.
40GBASE-FR PMD	1 024	2	25.6	Includes 2 m of fiber. See 89.3.1.
40GBASE-LR4 PMD	1 024	2	25.6	Includes 2 m of fiber. See 87.3.1.
40GBASE-ER4 PMD	1 024	2	25.6	Includes 2 m of fiber. See 87.3.1.
100G MAC, RS, and MAC Control	24 576	48	245.76	See 81.1.4.
100GBASE-R PCS	35 328	69	353.28	See 82.5.
100GBASE-R FEC	122 880	240	1 228.8	See 74.6.
100GBASE-R RS-FEC	40 960	80	409.60	See 91.4.
Inverse RS-FEC	40 960	80	409.60	See 152.4.
SC-FEC	1 827 840	3570	18 278.40	See 153.2.2.
100GBASE-R PMA	9 216	18	92.16	See 83.5.4.
100GBASE-P PMA	9 216	18	92.16	See 135.5.4.
100GBASE-KR2 PMD	4 096	8	40.96	Includes allocation of 20 ns for one direction through backplane medium. See 137.5.
100GBASE-KR4 PMD	2 048	4	20.48	Includes delay of one direction through backplane medium. See 93.4.
100GBASE-KP4 PMA/PMD	8 192	16	81.92	Includes delay of one direction through backplane medium. See 94.2.5.

100G

- 100G Module Delay
  - 100GBASE-R PMA: 92.16 ns / 4
  - Optical PMD: 20.48 ns
  - Total: 43.52 ns

Table 80–7—Sublayer delay constraints *(continued)*

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
100GBASE-CR2 PMD	4 096	8	40.96	Includes allocation for 20 ns for one direction through cable medium. See 136.5.
100GBASE-CR4 PMD	2 048	4	20.48	Does not include delay of cable medium. See 92.4.
100GBASE-CR10 PMD	9 728	19	97.28	Does not include delay of cable medium. See 85.4.
100GBASE-SR4 PMD	2 048	4	20.48	Includes 2 m of fiber. See 95.3.1.
100GBASE-SR2 PMD	2 048	4	20.48	Includes 2 m of fiber. See 138.3.
100GBASE-SR10 PMD	2 048	4	20.48	Includes 2 m of fiber. See 86.3.1.
100GBASE-DR PMD	2 048	4	20.48	Includes 2 m of fiber. See 140.3.
100GBASE-FR1 PMD	2 048	4	20.48	Includes 2 m of fiber. See 140.3.
100GBASE-LR4 PMD	2 048	4	20.48	Includes 2 m of fiber. See 88.3.1.
100GBASE-LR1 PMD	2 048	4	20.48	Includes 2 m of fiber. See 140.3.
100GBASE-ER4 PMD	2 048	4	20.48	Includes 2 m of fiber. See 88.3.1.
100GBASE-ZR PMD	2 048	4	20.48	Includes 2 m of fiber. See 154.3.1.

[a] For 40GBASE-R, 1 bit time (BT) is equal to 25 ps and for 100GBASE-R, 1 bit time (BT) is equal to 10 ps. (See 1.4.215 for the definition of bit time.)
[b] For 40GBASE-R, 1 pause_quantum is equal to 12.8 ns and for 100GBASE-R, 1 pause_quantum is equal to 5.12 ns. (See 31B.2 for the definition of pause_quanta.)
[c] Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.

100G Cont.

- 100G Module Delay
  - 100GBASE-P PMA: 92.16 ns / 3
  - Optical PMD: 20.48 ns
  - Total: 51.2 ns

Table 80–7—Sublayer delay constraints

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
...				
100GBASE-CR10 PMD	9 728	19	97.28	Does not include delay of cable medium. See 92.4.
100GBASE-VR1 PMD	2 048	4	20.48	Includes 2 m of fiber. See 167.3.
100GBASE-SR4 PMD	2 048	4	20.48	Includes 2 m of fiber. See 95.3.1.
100GBASE-SR2 PMD	2 048	4	20.48	Includes 2 m of fiber. See 138.3.
100GBASE-SR1 PMD	2 048	4	20.48	Includes 2 m of fiber. See 167.3.
100GBASE-SR10 PMD	2 048	4	20.48	Includes 2 m of fiber. See 86.3.1.
...				

[a] For 40GBASE-R, 1 bit time (BT) is equal to 25 ps and for 100GBASE-R, 1 bit time (BT) is equal to 10 ps. (See 1.4.215 for the definition of bit time.)
[b] For 40GBASE-R, 1 pause_quantum is equal to 12.8 ns and for 100GBASE-R, 1 pause_quantum is equal to 5.12 ns. (See 31B.2 for the definition of pause_quanta.)
[c] Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.

25G

- 25G Module Delay
  - PMA: 163.84 ns / 4
  - Optical PMD: 20.48 ns
  - Total: 61.44 ns

Table 105–3—Sublayer delay constraints

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
25G RS, MAC, and MAC Control	8 192	16	327.68	See 106.1.4.
25GBASE-R PCS	3 584	7	143.36	See 107.4.
25G BASE-R FEC	6 144	12	245.76	See 74.6.
25GBASE-R RS-FEC	24 576	48	983.04	See 108.4.
25GBASE-R PMA[d]	4 096	8	163.84	See 109.5.
25GBASE-CR PMD	512	1	20.48	See 110.4.
25GBASE-CR-S PMD	512	1	20.48	See 110.4.
25GBASE-KR PMD	512	1	20.48	See 111.4.
25GBASE-KR-S PMD	512	1	20.48	See 111.4.
25GBASE-SR PMD	512	1	20.48	See 112.3.
25GBASE-LR PMD	512	1	20.48	See 114.3.
25GBASE-ER PMD	512	1	20.48	See 114.3.
25GBASE-T PHY	25 600	50	1024	See 113.11.
25GBASE-BR10 PMD	512	1	20.48	See 158.3.
25GBASE-BR20 PMD	512	1	20.48	See 158.3.
25GBASE-BR40 PMD	512	1	20.48	See 158.3.

[a]1 bit time (BT) is equal to 40 ps. (See 1.4.215 for the definition of bit time.)
[b]1 pause_quantum is equal to 20.48 ns. (See 31B.2 for the definition of pause_quanta.)
[c]Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.
[d]Cumulative round-trip delay contributed by up to four PMA stages in a PHY.

200G

- 200G Module Delay
  - PMA: 92.16 ns / 4
  - Optical PMD: 20.48 ns
  - Total: 43.52 ns

Table 116–6—Sublayer delay constraints (200GBASE)

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
200G MAC, RS, and MAC Control	49 152	96	245.76	See 117.1.4.
200GBASE-R PCS or 200GXS[d]	160 256	313	801.28	See 119.5.
200GBASE-R PMA	18 432	36	92.16	See 120.5.4.
200GBASE-KR4 PMD	8 192	16	40.96	Includes allocation of 20 ns for one direction through backplane medium. See 137.5.
200GBASE-CR4 PMD	8 192	16	40.96	Includes allocation of 20 ns for one direction through cable medium. See 137.5.
200GBASE-SR4 PMD	4 096	8	20.48	See 138.3.
200GBASE-DR4 PMD	4 096	8	20.48	Includes 2 m of fiber. See 121.3.1.
200GBASE-FR4 PMD	4 096	8	20.48	Includes 2 m of fiber. See 122.3.1.
200GBASE-LR4 PMD	4 096	8	20.48	Includes 2 m of fiber. See 122.3.1.
200GBASE-ER4 PMD	4 096	8	20.48	Includes 2 m of fiber. See 122.3.1.

[a] For 200GBASE-R, 1 bit time (BT) is equal to 5 ps. (See 1.4.215 for the definition of bit time.)
[b] For 200GBASE-R, 1 pause_quantum is equal to 2.56 ns. (See 31B.2 for the definition of pause_quanta.)
[c] Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.
[d] If an implementation includes the 200GMII extender, the delay associated with the 200GMII extender includes two 200GXS sublayers.

400G

- 400G Module Delay
  - PMA: 92.16 ns / 4
  - Optical PMD: 20.48 ns
  - Total: 43.52 ns

Table 116–7—Sublayer delay constraints (400GBASE)

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
400G MAC, RS, and MAC Control	98 304	192	245.76	See 117.1.4.
400GBASE-R PCS or 400GXS[d]	320 000	625	800	See 119.5.
400GBASE-R PMA	36 864	72	92.16	See 120.5.4.
400GBASE-SR16 PMD	8 192	16	20.48	Includes 2 m of fiber. See 123.3.1.
400GBASE-SR8 PMD	8 192	16	20.48	Includes 2 m of fiber. See 138.3.1.
400GBASE-SR4.2 PMD	8 192	16	20.48	Includes 2 m of fiber. See 150.3.1.
400GBASE-DR4 PMD	8 192	16	20.48	Includes 2 m of fiber. See 124.3.1.

Table 116–7—Sublayer delay constraints (400GBASE) *(continued)*

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
400GBASE-FR8 PMD	8 192	16	20.48	Includes 2 m of fiber. See 122.3.1.
400GBASE-FR4 PMD	8 192	16	20.48	Includes 2 m of fiber. See 151.3.1.
400GBASE-LR4-6 PMD	8 192	16	20.48	Includes 2 m of fiber. See 151.3.1.
400GBASE-LR8 PMD	8 192	16	20.48	Includes 2 m of fiber. See 122.3.1.
400GBASE-ER8 PMD	8 192	16	20.48	Includes 2 m of fiber. See 122.3.1.

[a] For 400GBASE-R, 1 bit time (BT) is equal to 2.5 ps. (See 1.4.215 for the definition of bit time.)
[b] For 400GBASE-R, 1 pause_quantum is equal to 1.28 ns. (See 31B.2 for the definition of pause_quanta.)
[c] Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.
[d] If an implementation includes the 400GMII extender, the delay associated with the 400GMII extender includes two 400GXS sublayers.

maki_3dj_01_230427

50G

- 50G Module Delay
  - PMA: 92.16 ns / 3
  - Optical PMD: 20.48 ns
  - Total: 51.2 ns

Table 131–4—Sublayer delay constraints (50GBASE)

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
50G MAC, RS, and MAC Control	16 384	32	327.68	See 132.1.4.
50GBASE-R PCS	11 264	22	225.28	See 133.3.
50GBASE-R RS-FEC	25 600	50	512	See 134.4.
50GBASE-R PMA	4 608	9	92.16	See 135.5.4.
50GBASE-KR PMD	2 048	4	40.96	See 137.5. Includes allocation of 20 ns for the medium.
50GBASE-CR PMD	2 048	4	40.96	See 136.5. Includes allocation of 20 ns for the medium.
50GBASE-SR PMD	1 024	2	20.48	Includes 2 m of fiber. See 138.3.1.
50GBASE-FR PMD	1 024	2	20.48	Includes 2 m of fiber. See 139.3.
50GBASE-LR PMD	1 024	2	20.48	Includes 2 m of fiber. See 139.3.
50GBASE-ER PMD	1 024	2	20.48	Includes 2 m of fiber. See 139.3.
50GBASE-BR10 PMD	1 024	2	20.48	Includes 2 m of fiber. See 160.3.1.
50GBASE-BR20 PMD	1 024	2	20.48	Includes 2 m of fiber. See 160.3.1.
50GBASE-BR40 PMD	1 024	2	20.48	Includes 2 m of fiber. See 160.3.1.

[a]For 50GBASE-R, 1 bit time is equal to 20 ps. (See 1.4.215 for the definition of bit time.)
[b]For 50GBASE-R, 1 pause_quantum is equal to 10.24 ns. (See 31B.2 for the definition of pause_quanta.)
[c]Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.

800G

- 800G Module Delay
  - PMA: 92.16 ns / 4
  - Optical PMD: 20.48 ns
  - Total: 43.52 ns

169.3.3 Semantics of inter-sublayer service interface primitives

The semantics of the inter-sublayer service interface primitives for the 800GBASE-R sublayers are described in 116.3.3.1 through 116.3.3.3.

169.4 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementations conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices. Table 169–4 contains the values of maximum sublayer delay (sum of transmit and receive delays at one end of the link) in bit times as specified in 1.4 and pause_quanta as specified in 31B.2 for 800 Gigabit Ethernet. If a PHY contains an Auto-Negotiation sublayer, the delay of the Auto-Negotiation sublayer is included within the delay of the PMD and medium.

Table 169–4—Sublayer delay constraints (800GBASE)

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
800G MAC, RS, and MAC Control	196 608	384	245.76	See 170.1.4.
800GBASE-R PCS or 800GXS[d]	640 000	1250	800	See 172.5.
800GBASE-R PMA	73 728	144	92.16	See 173.4.4.
800GBASE-KR8 PMD	32 768	64	40.96	Includes allocation of 14 ns for one direction through backplane medium. See 163.5.
800GBASE-CR8 PMD	32 768	64	40.96	Includes allocation of 14 ns for one direction through cable medium. See 162.5.
800GBASE-VR8 PMD	16 384	32	20.48	Includes 2 m of fiber. See 167.3.1.
800GBASE-SR8 PMD	16 384	32	20.48	Includes 2 m of fiber. See 167.3.1.
800GBASE-DR8 PMD	16 384	32	20.48	Includes 2 m of fiber. See 124.3.1.
800GBASE-DR8-2 PMD	16 384	32	20.48	Includes 2 m of fiber. See 124.3.1.

[a] For 800GBASE-R, 1 bit time (BT) is equal to 1.25 ps. (See 1.4.215 for the definition of bit time.)
[b] For 800GBASE-R, 1 pause_quantum is equal to 640 ps. (See 31B.2 for the definition of pause_quanta.)
[c] Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.
[d] If an implementation includes the 800GMII Extender, the delay associated with the 800GMII Extender includes two 800GXS sublayers.

100GBASE-ZR

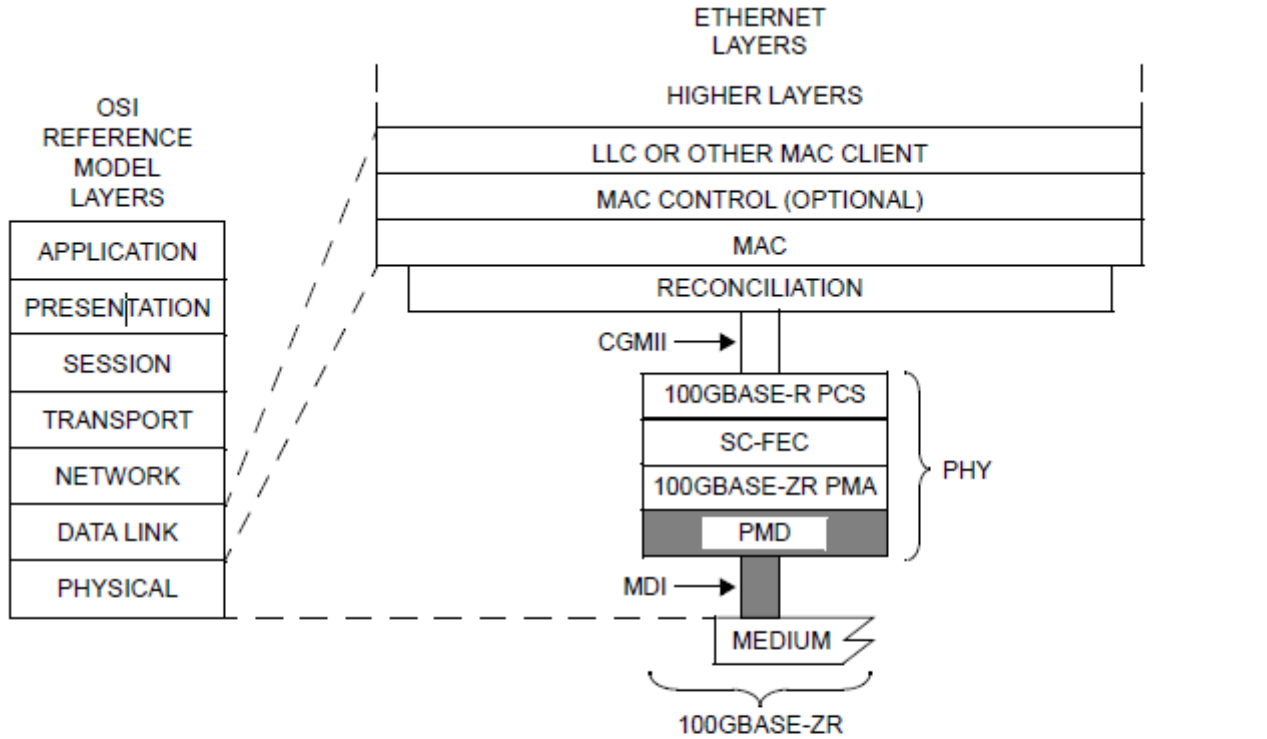

Figure 154–1—100GBASE-ZR PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

154.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 100GBASE-ZR PMD including 2 m of fiber in one direction shall be no more than 2048 bit times (4 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

100GBASE-ZR

- 100GBASE-ZR Module Delay
  - PMA: 92.16 ns / 3
  - Inverse RS-FEC: 409.6 ns
  - SC-FEC: 18,278.40 ns
  - 100GBASE-ZR PMA: ?
  - PMD: 20.48 ns
  - Total: ?

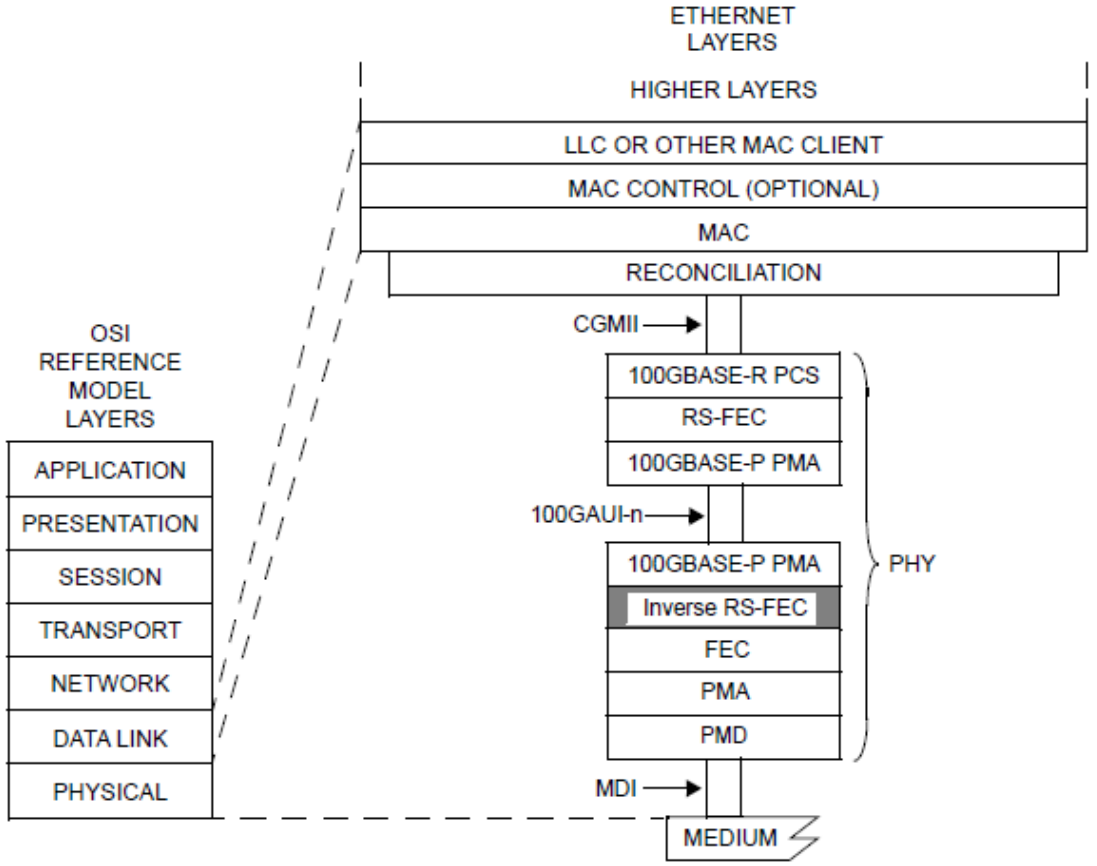

100GAUI-n = 100 Gb/s n-LANE ATTACHMENT UNIT INTERFACE
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
Inverse RS-FEC = Inverse REED-SOLOMON FORWARD ERROR CORRECTION
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

Figure 152–1—Inverse RS-FEC relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

152.4 Delay constraints

The maximum delay contributed by the Inverse RS-FEC sublayer (sum of transmit and receive delays at one end of the link) shall be no more than 40 960 bit times (80 pause_quanta or 409.6 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

400GBASE-ZR

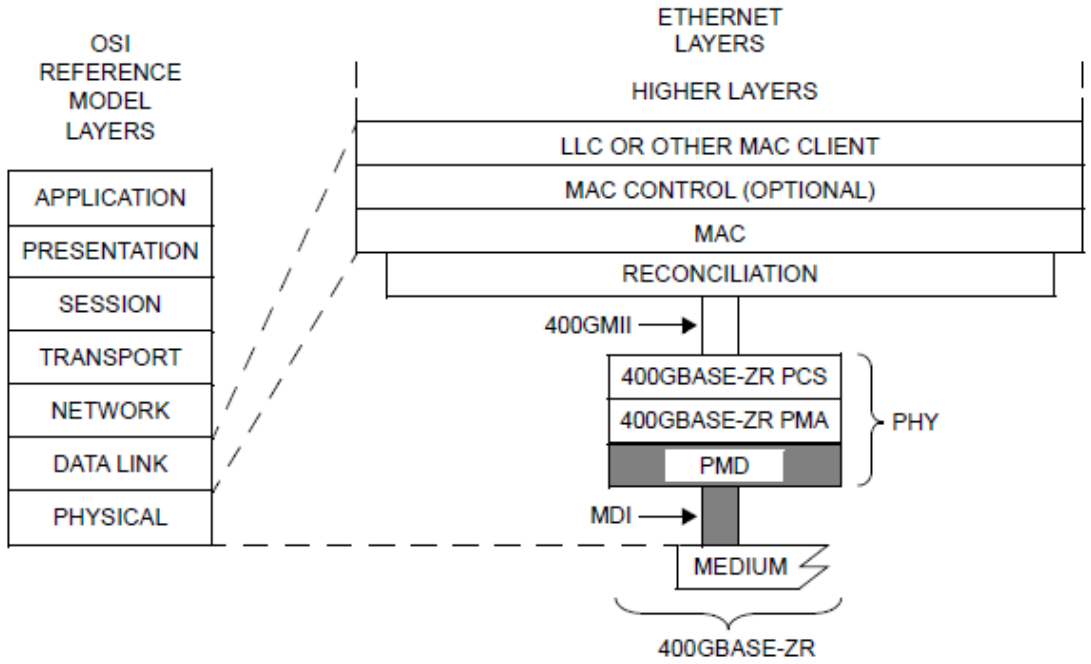

400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
400GBASE-ZR = PMD FOR DWDM CHANNEL OVER A DWDM BLACK LINK

Figure 156–1—400GBASE-ZR PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

Replace Figure 118–1 with the following figure:

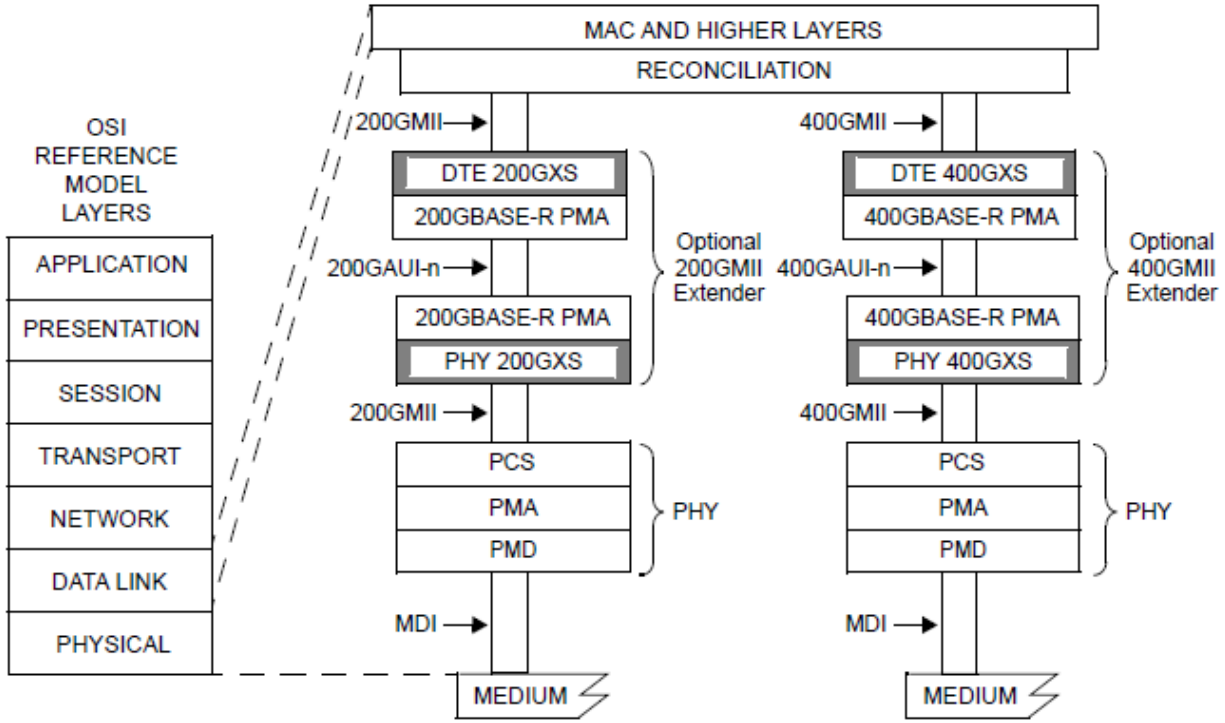

200GAUI-n = 200 Gb/s n-LANE ATTACHMENT UNIT INTERFACE
200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
200GXS = 200GMII EXTENDER SUBLAYER
400GAUI-n = 400 Gb/s n-LANE ATTACHMENT UNIT INTERFACE
400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
400GXS = 400GMII EXTENDER SUBLAYER
DTE = DATA TERMINAL EQUIPMENT
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

Figure 118–1—200GXS and 400GXS relationship to the ISO/IEC Open System Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

maki_3dj_01_230427

400GBASE-ZR

- 400GBASE-ZR Module Delay
  - 400GBASE-R PMA: 92.16 ns / 4
  - 400GXS: 800 ns
  - 400GBASE-ZR PCS and 400GBASE-ZR PMA: 6,000.64 ns
  - Total: 6,823.68 ns

Table 116–7—Sublayer delay constraints (400GBASE)

Sublayer	Maximum (bit time)[a]	Maximum (pause_quanta)[b]	Maximum (ns)	Notes[c]
…				
400GBASE-ZR PCS and PMA	2 400 256	4688	6000.64	
400GBASE-ZR PMD	8192	16	20.48	Includes 2 m of fiber. See 156.3.

[a] For 400GBASE-~~R~~, 1 bit time (BT) is equal to 2.5 ps. (See 1.4.215 for the definition of bit time.)
[b] For 400GBASE-~~R~~, 1 pause_quantum is equal to 1.28 ns. (See 31B.2 for the definition of pause_quanta.)
[c] Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.

maki_3dj_01_230427

Summary

- For the PAUSE application, the key parameter is the total Delay (overall system latency)
- Delay is the sum of the Tx and Rx paths
- Optical module portion of the Delay depends upon the PMA and PMD Sublayers in all cases
  - PMD Sublayer in the module
  - PMA Sublayer distributed across the module and host
- In some cases, the module has Delay from PCS/FEC
- The whole link includes 3 or 4 stages of the PMA Sublayer
  - Optical module typically has 1 stage but sometimes has more