

IEEE 802.3 Ethernet Working Group
DRAFT Liaison Communication

Source: IEEE 802.3 Working Group¹

To: Glenn Parsons Chair, ITU-T SG15
[REDACTED]
Hiroshi Ota Advisor, ITU-T SG15
[REDACTED]
Steve Gorshe Rapporteur Q11/15
[REDACTED]

CC: Alpesh Shah Secretary, IEEE-SA Standards Board
Secretary, IEEE-SA Board of Governors
[REDACTED]
James Gilb Chair, IEEE 802 LMSC
[REDACTED]
Adam Healey Vice-chair, IEEE 802.3 Ethernet Working Group
[REDACTED]
Jon Lewis Secretary, IEEE 802.3 Ethernet Working Group
[REDACTED]
John D'Ambrosia Chair, IEEE P802.3dj Task Force
[REDACTED]
Mark Nowell Vice-Chair, IEEE P802.3dj Task Force
[REDACTED]

From: David Law Chair, IEEE 802.3 Ethernet Working Group
[REDACTED]

Subject: Liaison to ITU-T SG15 requesting FlexO payload types for 800GBASE-ER1

Approval: Agreed to at IEEE 802.3 interim meeting, Hamburg, 19 September 2024

Dear Mr. Parsons and members of ITU-T SG15,

We would like to make you aware that some of our work on 800G coherent PHYs in the IEEE P802.3dj project is using the FlexO-8e frame structure and FlexO-8e-DO FEC frame structure. The 800GBASE-ER1 PHY uses those frame structures with an optical interface intended to support 40 km applications, while the 800GBASE-ER1-20 PHY uses them with an optical interface intended to support 20 km applications. Both PHYs are single-channel (non-DWDM) interfaces using a single wavelength in the C band. We also intend that these two PHYs will be interoperable when connected to each other over a channel that is compliant with the channel specified for 800GBASE-ER1-20.

¹ This document solely represents the views of the IEEE 802.3 Working Group, and does not necessarily represent a position of the IEEE, the IEEE Standards Association, or IEEE 802.

When we started our work on these PHYs, we intended to align fully with the 800ZR frame that OIF is in the process of specifying (which is also based on the FlexO-8e and FlexO-8e-DO frames) and to use the same FlexO payload types that OIF uses in the 800ZR application. In the course of our work, we have found a need to improve the accuracy of PTP over the 800GBASE-ER1 and 800GBASE-ER1-20 interfaces to satisfy the needs of some users of these interfaces. The inaccuracy is caused by the removal of 800GBASE-R alignment markers prior to mapping into the FlexO-8e frame and re-insertion of those alignment markers, potentially in a different location in the client bitstream, when de-mapping from the FlexO-8e frame. A change in location of the alignment markers can cause a PTP packet that is within the client data stream to move forward or backward in time and introduce a timestamp error of ~5.12 ns at 800 Gb/s rates. To mitigate this error, we have chosen to define some additional mapping-specific overhead in reserved bytes of the FlexO-8e frame. This overhead is used to convey information about where the alignment markers are located such that the far end is able to re-insert them in the same position from where they were removed from the client bitstream. The location of the additional overhead is indicated in Figure 1.

For additional detail about how the overhead is used, please refer to this presentation: https://www.ieee802.org/3/dj/public/24_05/sluyski_3dj_01a_2405.pdf. Note that we are considering alternative names for these new overhead bytes, as they are not really related to the justification control information that is present in JC1-JC6.

Frame	MFAS bits	11- 13- 29-											
		1	2	3	4	5	6	7-10	12	26	27-28	40	
1	xxxxx000	MFAS	STAT	GID	GID	GID	RES	IID	MAP	CRC	FCC1	OSMC	RES
2	xxxxx001	MFAS	STAT	AVAIL	RES	JC4	JC1	MAP	CRC	FCC2	OSMC	RES	
3	xxxxx010	MFAS	STAT	JC7	JC8	JC5	JC2	MAP	CRC	FCC3	OSMC	RES	
4	xxxxx011	MFAS	STAT	RES	JC9	JC6	JC3	MAP	CRC	FCC4	OSMC	RES	
5	xxxxx100	MFAS	STAT	RES	RES	MSI	PT	MAP	CRC	FCC5	OSMC	RES	
6	xxxxx101	MFAS	STAT	RES	RES	JC4	JC1	MAP	CRC	FCC6	OSMC	RES	
7	xxxxx110	MFAS	STAT	JC7	JC8	JC5	JC2	MAP	CRC	FCC7	OSMC	RES	
8	xxxxx111	MFAS	STAT	RES	JC9	JC6	JC3	MAP	CRC	FCC8	OSMC	RES	

Additional mapping-specific overhead used to improve PTP accuracy in 800GBASE-ER1 applications
 FlexO overhead not used in 800GBASE-ER1 applications

Figure 1: New mapping-specific overhead for improving PTP accuracy

With the addition of this new overhead to the frame, interoperation cannot be guaranteed if an 800GBASE-ER1 interface and an 800ZR interface are connected, and therefore it is no longer appropriate for the 800GBASE-ER1 and 800GBASE-ER1-20 interfaces to use the same FlexO payload types as OIF 800ZR. As such, we request that you please assign a FlexO payload type for use by 800GBASE-ER1 and 800GBASE-ER1-20 applications.

The next meetings of the IEEE P802.3dj task force are the weeks of 11 November 2024, 20 January 2025, and 10 March 2025. We very much appreciate the ability to work collaboratively with ITU-T SG15 on this and other matters of common interest.

Sincerely,

David Law
 Chair, IEEE 802.3 Ethernet Working Group