400G TECHNICAL POSTER

Coherent optics

Digital coherent optics (DCO)

- At its most basic, this technology utilizes low-cost standard optics transmitting 400G Ethernet over very long distances using DWDM and higher modulation schemes such as Quadrature Amplitude Modulation (e.g., 8QAM and 16QAM)
- Coherent optics are available in QSFP-DD, OSFP or CFP2 form factors—and are replacing coherent line cards in the field
- Uses amplitude and phase modulation of light
- Uses dual polarization
- Transports considerably more information through optical cables

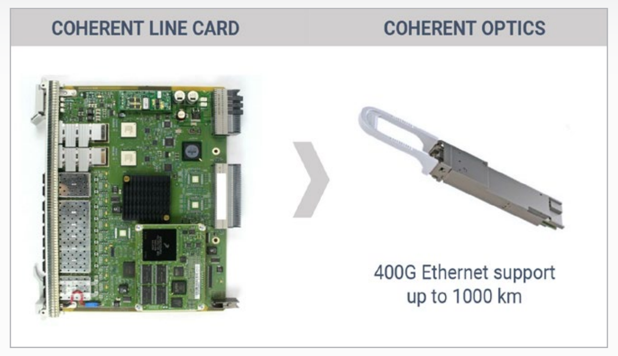

Why?

The industry needs standard DWDM technology in a smaller footprint to transport various clients over longer distances.

Conceptual approach

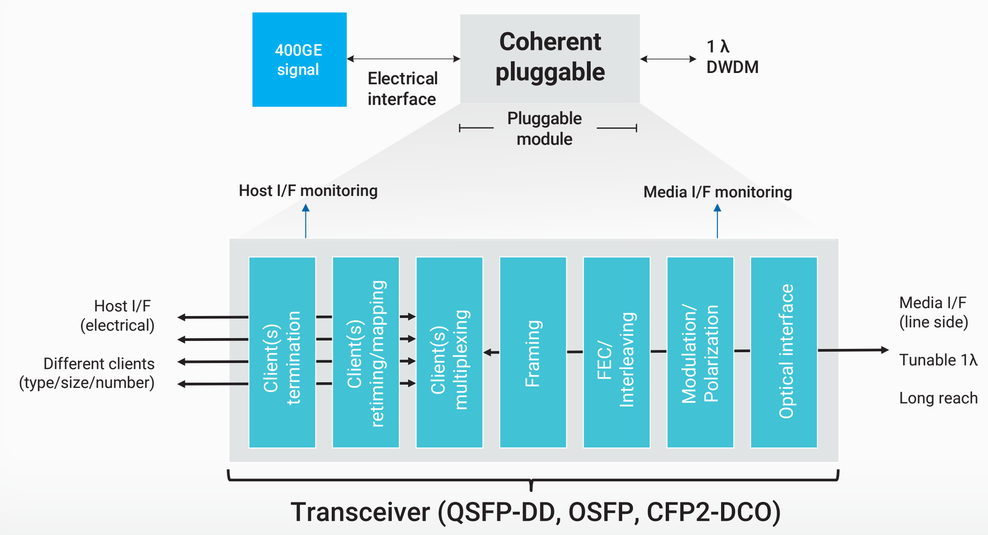

Standards

- Several DCO standards exist (e.g., OIF 400ZR, IEEE 802.3cw, OpenZR+, OpenROADM)
- These DCO standards share similarities while bearing noteworthy differences
- The table below summarizes these differences

Optical interfaces	400ZR	OpenZR+	OpenROADM
Line rate / Capacity	400G	100/200/300/400G	100/200/300/400G
Client types	400GE	100/200/400GE	100/200/400GE OTU4, OTUC1-4
Client multiplexing	No	Yes No mix	Yes Mixed types
Modulation	400G: DP-16QAM	400G: DP-16QAM 300G: DP-8QAM 200G: DP-QPSK 100G: DP-QPSK	400G: DP-16QAM 300G: DP-8QAM 200G: DP-16QAM 200G: DP-QPSK 100G: DP-QPSK
FEC	CFEC	OFEC	OFEC
Reach	120 km	>1000 km	1000 km

Standards

OIF 400G ZR

The OIF 400 ZR protocol aims to transparently transport a 400GBASE-R signal:

- Over a distance up to 120 km
- On a single DWDM wavelength
- On a point-to-point link (black link)
- Using a pluggable module
- With strong error correction capability (BER) of 1 x 10-15 out when submitted to an input BER of 1.25 x 10-2)

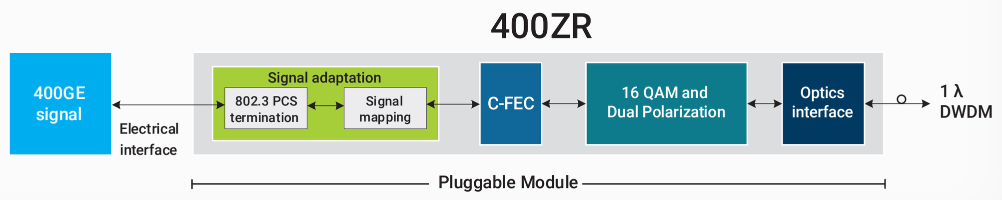

OpenZR+

The OpenZR+ protocol aims to transparently transport 100G, 200G or 400G Ethernet client signals:

- Over hundreds of kilometers on a single DWDM wavelength
- Using QPSK, 8QAM, and 16QAM optical modulation and dual polarization
- Using strong error correction capability with open forward error correction (OFEC)
- Using a pluggable module

Key benefits of OpenZR+ specification compared to 400ZR are:

- Client multiplexing modes
- Longhaul performance
- Extended reach for black link applications

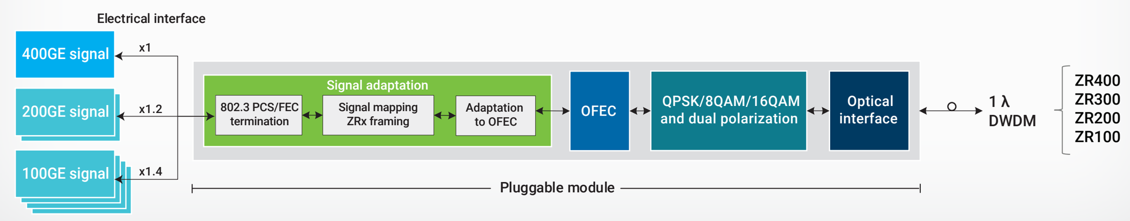

OpenROADM

- OpenROADM defines a protocol that can transparently transport 100/200/400G Ethernet, OTUC1-4, and OTU4 client(s) signals:
- On a single DWDM wavelength with a reach of up to 1000 km
- Using QPSK, 8QAM, and 16QAM optical modulation and dual polarization (rate dependent)
- Using strong error correction capability with OFEC
- Using a pluggable module

- The goals of the Open ROADM Multi-Source Agreement are:
  1. The disaggregation and opening of traditionally proprietary ROADM systems
  2. The SDN enablement of traditionally fixed ROADMs
- There are 2 specification series:
  1. Single-Wave (or W) to define how pluggable optics or transponders interoperate
  2. Multi-Wave (or MW) to define how ROADMs interoperate

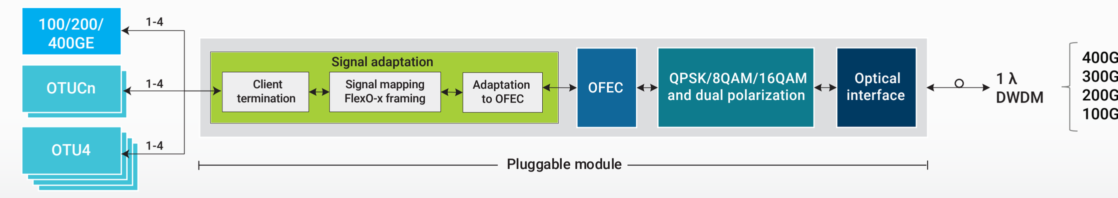

Module blocks

Signal adaptation

- The Ethernet signal from the electrical interface is processed to terminate the IEEE 802.3 PCS and RS-FEC
- For 200GE and 400GE clients, the FEC degrade condition is propagated and mapped to the ZR frame overhead
- Rate adaptation between client side and the line side is done using the Generic Mapping Procedure (GMP) defined in ITU-T G.709
- The ZR400 frame is same as OIF 400ZR frame (adapted FlexO frame). ZR300/200/100 frames vary in size.
- Multiplexing of Ethernet clients into the ZRx frame is performed using simple time division multiplexing (100G timeslots)
- The ZRx frame structure is adapted to the OFEC coder block by the addition of padding bits to create the required format of OFEC input block

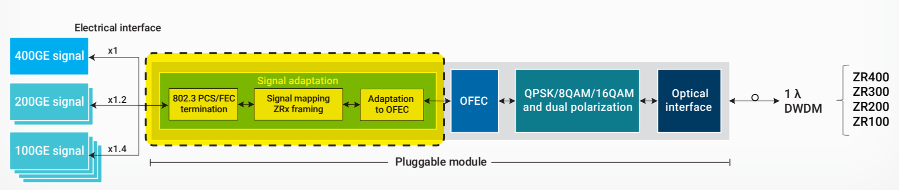

Forward error correction (FEC)

The FEC used depends on the standard. OpenZR+ and OpenROADM use OFEC:

- The OFEC encoding block consists of two FEC encoders/decoders engines operating in parallel
- The OFEC engine is a block-based encoder and iterative soft-decision decoder
- Uses the extended BCH (256,239) code

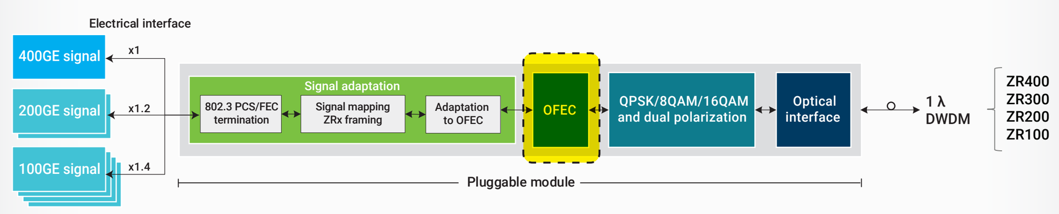

The following table shows the performance of the OFEC compared to the CFEC used by OIF 400ZR

FEC	OH	Latency	BER in at BER out = 1E-15	Net coding gain (16QAM)
OFEC	15.3%	~0.4 us	26.5625 Gbit/s	11.6 dB
CFEC	14.8%	~2.3 us	53.125 Gbit/s	10.8 dB

Modulation and dual polarization

The bits read out of the OFEC are used to form symbols in the X polarization (even bits) and in the Y polarization (odd bits).

Data bits are mapped to I and Q values using Gray coding (modulation dependent).

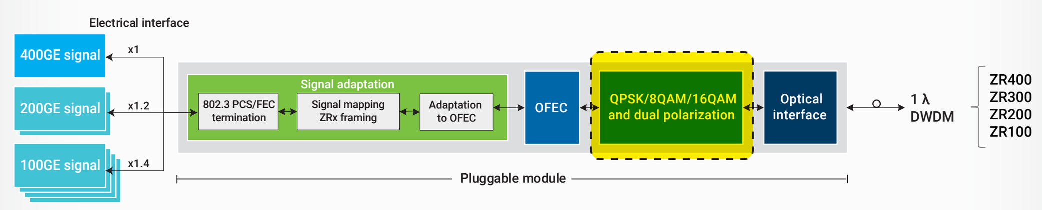

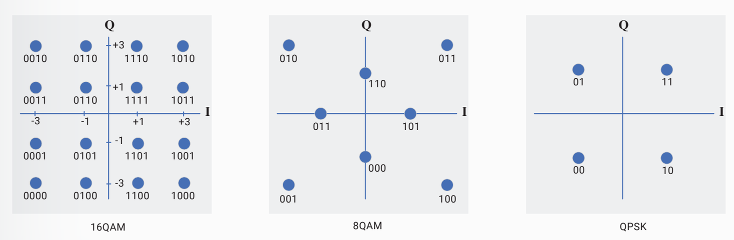

A frame structure (DSP Super-frame) is defined in the "Symbols Domain" to manage the signal recovery.

Overhead symbols are added to help the receiver to resolve the structure of the received encoded data.

Optical interface

Translates digital information carried by X and Y polarization paths into an analog coherent optical signal.

The signal processing is as follows:

- QPSK, 8QAM, or 16QAM modulation on each polarization axis using the I and Q data
- Modulation done over selected DWDM wavelength
- Combination of X and Y branches into one optical signal

The signal out of the optical interface complies to the optical parameters defined in the specific standard.

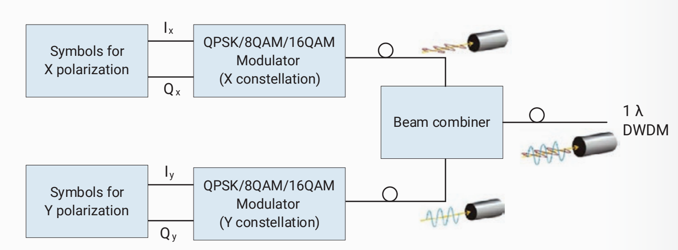

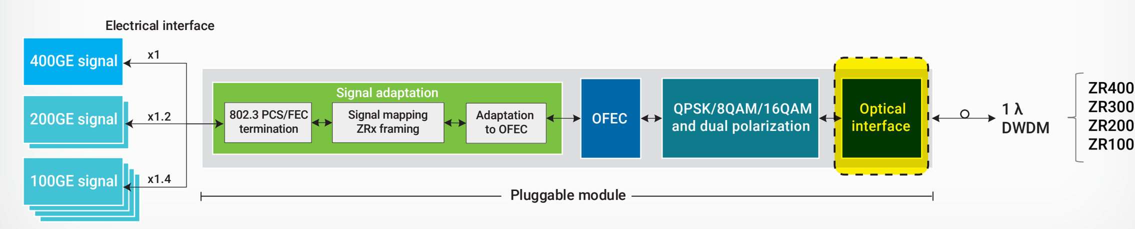

Coherent interface

Media interfaces

The diagram shows the various media interfaces supported by different standards.

- FOICx.k-DO meaning:
  - x = FlexO-x frame structure: 1 = FlexO-1 (100G), 2 = FlexO-2 (200G), 3 = FlexO-3 (300G), 4 = FlexO-4 (400G)
  - k = modulation order: 4 = QPSK, 6 = 8QAM, 8 = 16QAM
  - DO = protected with OFEC
- Multiplexing of clients supported (e.g., 4x100GE, 4xOTU4)
- Mix of clients allowed (e.g., 1x100GE + 1xOTUC2 + 1xOTU4)

Standard	Media interface	Payload capacity	FEC	Modulation	Symbol baud rate (per polarization)	Bit rate	Client
OIF 400GZR	1 x 400GE	400G	CFEC	16QAM	59 843 750 000	478 750 000 000	400GE
OpenZR+ & OpenROADM	1 x 400GE 2 x 200GE 4 x 100GE	400G	OFEC	16QAM	60 138 546 798	481 108 374 384	100GE
OpenZR+ & OpenROADM	3x100GE	300G	OFEC	8QAM	60 138 546 798	360 831 280 788	100/200GE
OpenZR+ & OpenROADM	1 x 200GE 2 x 100GE	200G	OFEC	QPSK	60 138 546 798	240 554 187 192	100/200GE
OpenZR+ & OpenROADM	1 x 100GE	100G	OFEC	QPSK	30 069 273 399	120 277 093 596	100GE
OpenROADM	FOIC4.8-DO (16QAM)	400G	OFEC	16QAM	63 139 467 923	505 115 743 384	100/200/400GE, OTUC1,2,3,4, OTU4
OpenROADM	FOIC3.6-DO (8QAM)	300G	OFEC	8QAM	63 139 467 923	378 836 807 538	100/200GE, OTUC1,2,3, OTU4
OpenROADM	FOIC2.4-DO (QPSK)	200G	OFEC	QPSK	63 139 467 923	252 557 871 692	100/200GE, OTUC1,2, OTU4
OpenROADM	FOIC2.8-DO (16QAM)	200G	OFEC	16QAM	31 569 733 961	252 557 871 688	100/200GE, OTUC1,2, OTU4
OpenROADM	FOIC1.4-DO (QPSK)	100G	OFEC	QPSK	31 569 733 961	126 278 935 844	100GE, OTUC1, OTU4

Modulation schemes

Coherent pluggable rely on higher order modulation to transport different clients over longer distances. Depending on the client transported by the pluggable, different orders of quadrature amplitude modulation (QAM) may be used.

The advantage: these schemes can carry more bits of information per symbol (e.g., 16QAM carries 4 bits per symbol).

The constellation diagrams below show data bits mapped to I and Q values using Gray coding for various forms of QAM.

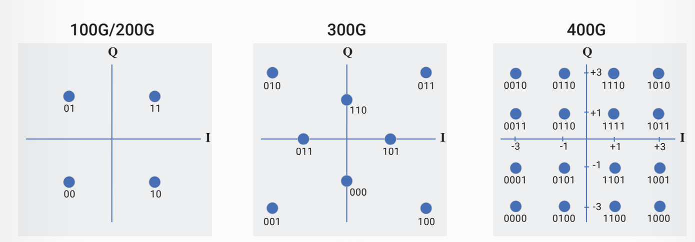

Interfaces

CFP2-DCO

- Targeted for metro/core applications for coherent 100G and 200G
- MDIO support

QSFP-DD

- High port density
- Accommodates full range of rates
- I2C and CMIS support

OSFP

- High thermal capability
- Accommodates full range of optics
- I2C and CMIS support

EXFO solutions

Test in the lab, field or on the production floor

Open Transceiver System (OTS): supports today's and tomorrow's optical interfaces

EXFO

400G TECHNICAL POSTER

400G/200G Ethernet

Ethernet frame format and rates

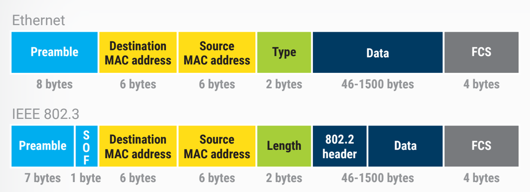

Ethernet interface	Line rate
400G Ethernet	425 Gbit/s

Packetized data

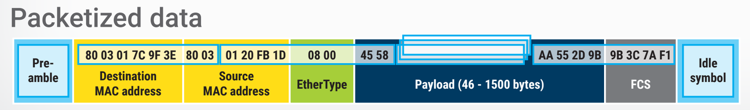

IEEE 802.3bs highlights

- Support a MAC data rate of 200 Gbit/s
- Support a MAC data rate of 400 Gbit/s
- Use of FEC is mandatory (544,514)
- Preserve minimum and maximum frame size
- Include Energy Efficient Ethernet (EEE)
- Provide BER < 10-13 at MAC layer

400GE packet transmission

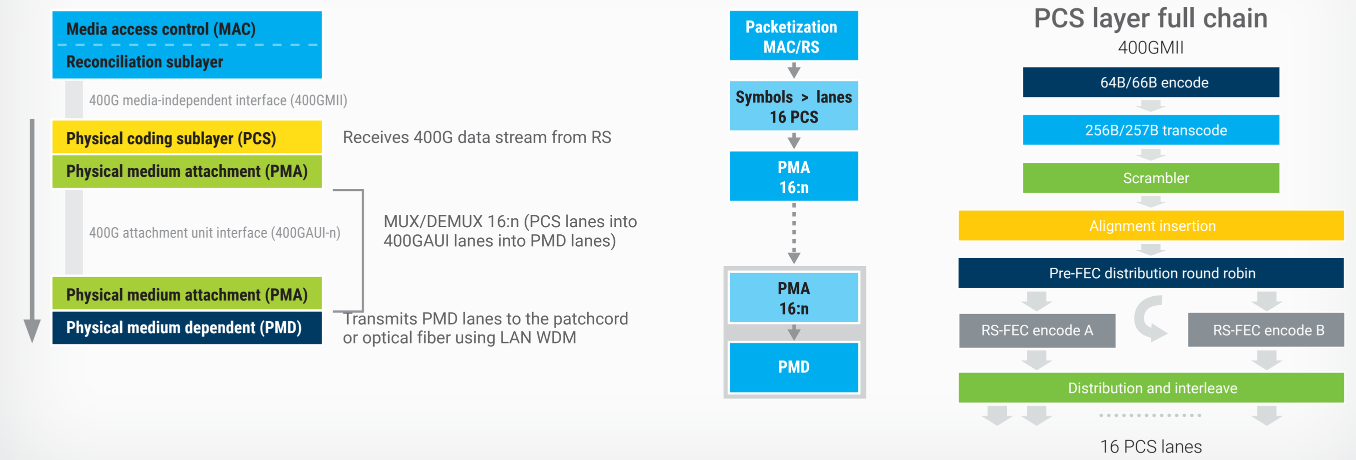

PCS lane skew

- Skew is the difference between the times of the earliest PCS lane and latest PCS lane for the 1 to 0 transition of the alignment marker sync bits.
- Skew variation may be introduced due to electrical, thermal or environmental variations.

Skew points	Maximum skew (ns) [a]	Maximum skew for 200GBASE-R or 400GBASE-R PCS lane (UI) [b]
SP1	29	≈ 770
SP2	43	≈ 1142
SP3	54	≈ 1434
SP4	134	≈ 3559
SP5	145	≈ 3852
SP6	160	≈ 4250
At PCS receive	180	≈ 4781

a. Skew limit includes 1 ns allowance for PCB traces that are associated with the skew points.

b. The symbol ≈ indicates approximate equivalent of maximum skew in UI, based on 1 UI equals 37.64706 ps at PCS lane signaling rate of 26.5625 GBd.

Maximum skew and skew variation at physically instantiated interfaces are specified at skew points SP1, SP2 and SP3 for the transmit direction and SP4, SP5 and SP6 for the receive direction.

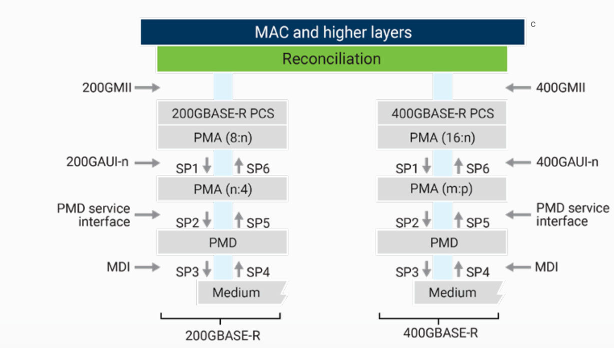

c. Source: 802.3bs-2017, IEEE Standard for Ethernet

Legend	
200GAUI-n	200 Gbit/s attachment unit interface
200GMII	200 Gbit/s media-independent interface
400GAUI-n	200 Gbit/s media-independent interface
400GMII	400 Gbit/s media-independent interface
MAC	Media access control
MDI	Medium-dependent interface
PCS	Physical coding sublayer
PMA	Physical medium attachement
PMD	Physical medium dependent
n	8 or 4
m	16 or 8
p	16, 8 or 4

Flex Ethernet (FlexE)

FlexE capabilities

- FlexE Version 1.0 is an implementation agreement published by the Optical Interworking Forum (OIF)
- It is a generic mechanism that enables the support of a multitude of Ethernet MAC rates that may or may not correspond to an actual Ethernet PHY rate
- It provides a flexible physical mapping structure to transport different Ethernet rates
- It is the alternative to other aggregation protocols in layer 3
- FlexE provides easy integration for future Ethernet rates

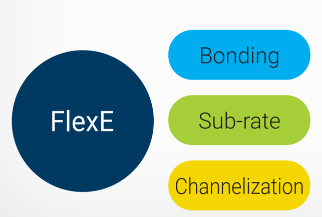

- Creates high capacity data pipes by aggregating several Ethernet PHYs to transport MAC rates greater than 100G
- An action that allows network elements to sub-divide physical interfaces in order to transport lower data pipes over partially filled Ethernet PHYs
- Provides network elements with the capability to create specific transport channels for multiple data pipes through associated Ethernet PHYs, on the same or even different directions

FlexE general structure

The FlexE group is composed from 1 to n bonded Ethernet PHYs.

- The version 1.0 of FlexE supports only 100GBASE-R PHYs
- Higher rates are intended to be added once available

The FlexE client is an Ethernet flow based on a MAC data rate that may or may not correspond to an existing Ethernet PHY rate.

- The FlexE client MAC rates supported are 10, 40 and $n \times 25$G

The FlexE shim is the layer that maps and demaps the FlexE clients transported over a FlexE group.

- The MUX operates in the transmit direction
- The DEMUX operates in the receive direction

FlexE shim layer

The FlexE shim maps or demaps the FlexE clients carried over a FlexE group. The FlexE shim layer is located in the middle of the PCS in the 100GBASE-R.

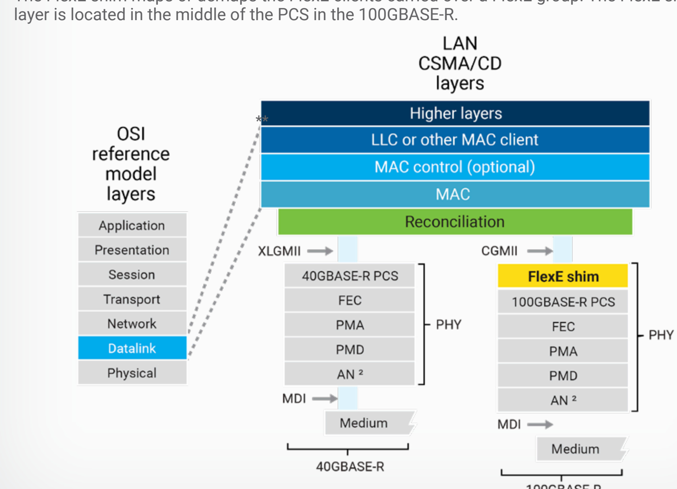

FlexE calendar

- FlexE uses a calendar that assigns 66B block positions on sub-calendars for each PHY that belongs to the group, for each FlexE client. The calendar granularity is 5G, and the length is 20 slots per 100G of group capacity.
- Each client is assigned to a particular place in the master calendar, somewhat similar to how it is done with transport, where a tributary is assigned to a higher-order frame. Clients are inserted in the master calendar in terms of 64/66 blocks.

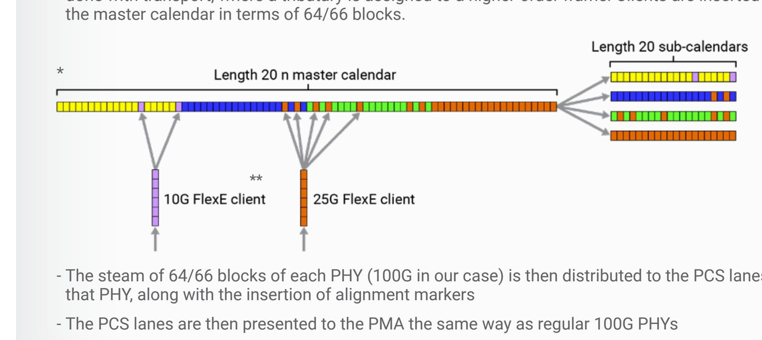

- The steam of 64/66 blocks of each PHY (100G in our case) is then distributed to the PCS lanes of that PHY, along with the insertion of alignment markers
- The PCS lanes are then presented to the PMA the same way as regular 100G PHYs
- If 802.3bj FEC is used, these 64/66 blocks are transcoded to 256/257 and FEC encoded

* Source: OIF-FLEXE-01.0, Flex Ethernet, Implementation Agreement, Optical, Internetworking Forum (OIF)

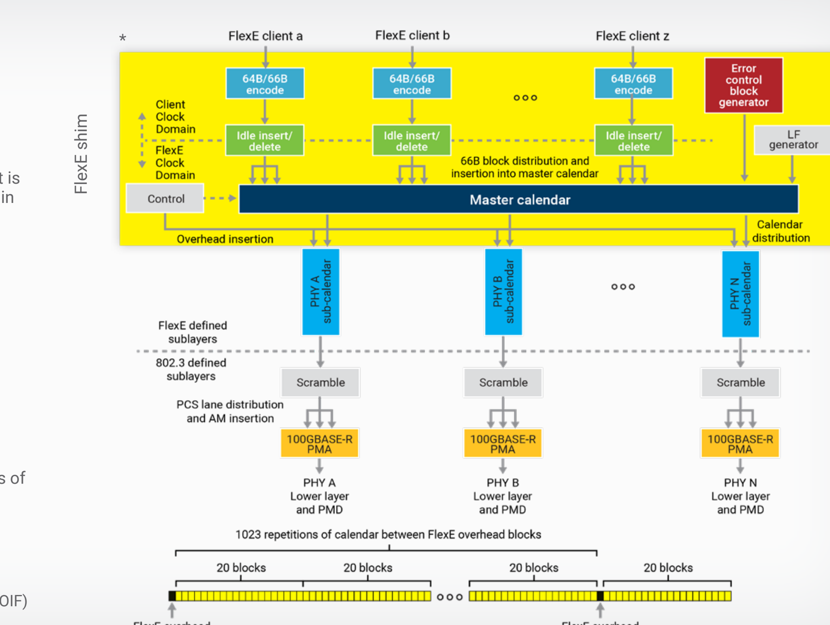

FlexE BERT

Our FTBx-88460NGE Power Blazer 400G multiservice tester is the most compact solution on the market and includes basic and advance capabilities for lab and field implementations, including our Flex BERT application.

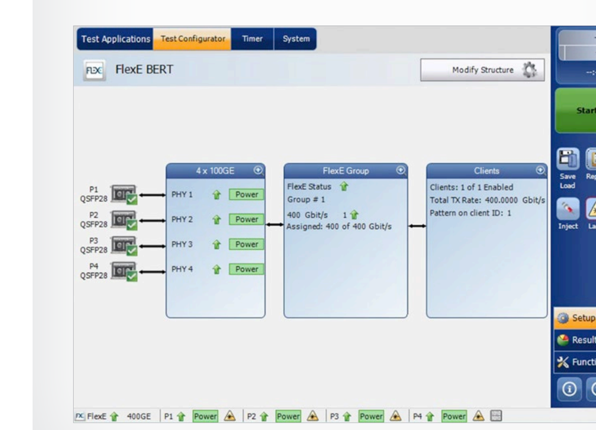

OTUCn/FlexO

OTN evolution

- ITU G.709 was adapted in 2016 with the addition of Y.1331 “OTUCn”
- OTUCn, the “C” corresponds to the roman numeral for 100
- OTUCn associates PHYs thru FlexO (almost like FlexE)
- The objective of OTN beyond 100G technology is to provide a long term evolution of an OTN protocol that can carry, in a flexible manner, payloads demanding bandwidth above 100 Gbit/s—such as 200GE, 400GE and FlexE

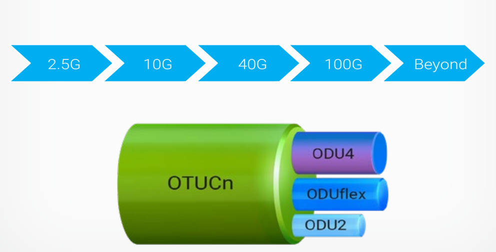

OTUCn highlights

The OTUCn signal is made of the same classical OTN constituents but adapted to handle a concept of “slice interleaving” similar to SONET/SDH that is used to increase the transport capacity.

OTUCn functionality involves the following:

- Client mapping in ODUk
- Mapping ODUk containers in the OPUCn payload area
- Overhead generation and extraction

A basic frame structure with 100 Gbit/s capacity is defined and referred to as OTUC.

The payload area is a contiguous container made of tributary slots, these are part of the OPUCn.

FEC is not part of the basic frame since there will be different interface types (e.g., 100G, 200G, SR, LR) to carry the OTUCn and each interface will have its own requirements in terms of strength of FEC required.

There are 20 tributary slots entities per OTUC slice:

- Tributary slot rate: 5.24089 Gbit/s, granularity 16 bytes

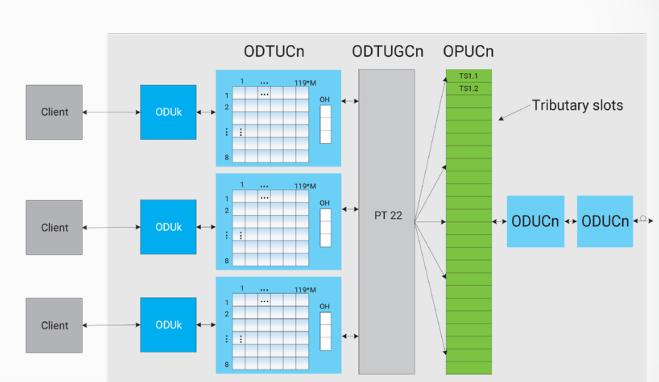

OTUCn frame structure

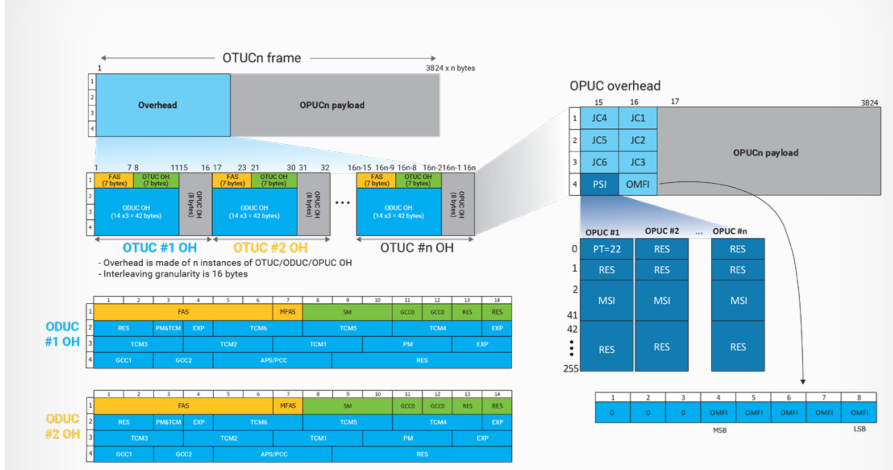

Flexible OTN (FlexO) highlights

FlexO is an adaptation layer which provides a flexible, modular mechanism to support different line rates beyond 100G signals.

FlexO enables a set of n 100 Gbit/s PHYs to be bound together to carry an OTUCn, with each 100 Gbit/s PHY carrying an OTUC slice.

FlexO provides frame alignment, deskew, group management and a management communication channel.

FlexO processes an OTUCn signal as follows:
- Source: splits the OTUCn frame into n * OTUC slices
- Sink: combines n * OTUC instances into an OTUCn

FlexO functional block

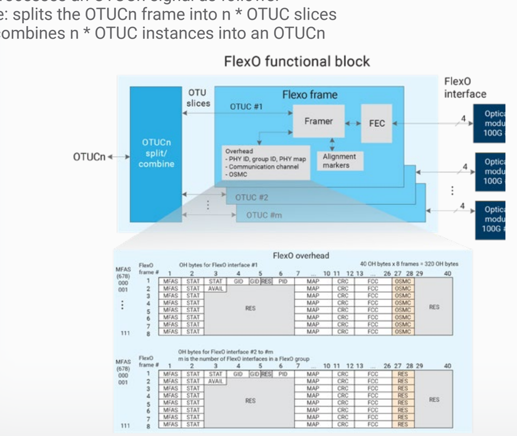

OTUCn/FlexO signal flow

Figure below includes all structures together and depicts signal flow from the client signal to the optical module.

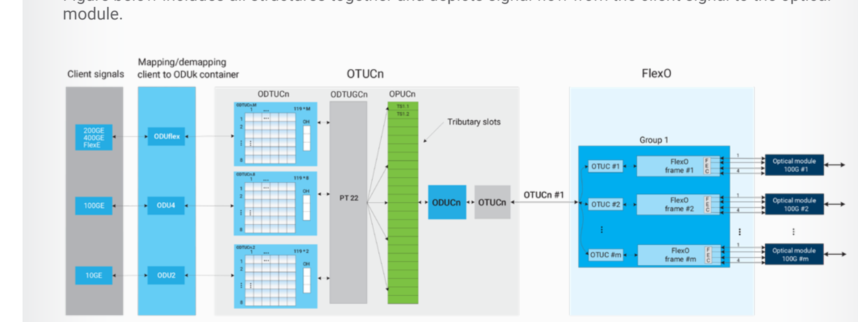

FlexO BERT

EXFO’s FTBx-88460 multiservice test solution supports advanced capabilities. Ready for lab and field implementations, it covers the complete 400G ecosystem, and includes our FlexO BERT application.

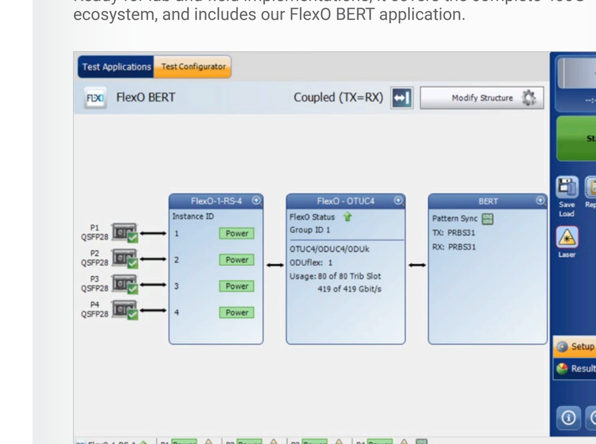

400G/200G interfaces and pluggable transceivers

Physical-layer specifications

IEEE 802.3bs standard provides physical-layer specifications that support:

Physical characteristics	Optical interfaces	PCS lanes	Optical lanes	Modulation	Rate
100 m over MMF	400GBASE-SR16	16	16	NRZ	26.5625 Gbit/s
2 km over SMF	400GBASE-FR8	16	8	PAM4	53.125 Gbit/s
10 km over SMF	400GBASE-LR8	16	8	PAM4	53.125 Gbit/s
500 m over MMF	400GBASE-DR4	16	4	PAM4	106.25 Gbit/s
2 km over SMF	400GBASE-FR4	16	4	PAM4	106.25 Gbit/s

New modulation schemes

New modulation schemes allow the transporting of more data in the same transceiver form factors.

PAM4 signaling delivers twice as many bits per sample

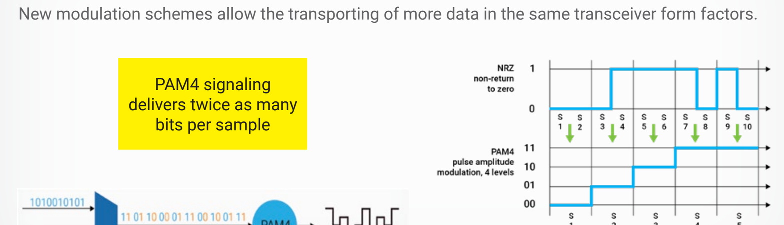

PAM4 challenges

The complexity of the new modulation (PAM4) relating to 400G technologies creates new testing challenges for labs. Technicians require a single tester to be able to effectively manipulate the transmission signal and receiving signal.

Solutions need to enable users in handling and monitoring the following parameters:

- Pre-Emphasis
- RX- Equalization
- PAM4 Histogram

Interfaces

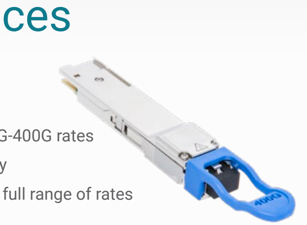

QSFP-DD
- Support for 200G-400G rates
- High port density
- Accommodates full range of rates

OSFP
- Supports 400G and 800G
- High thermal capability
- Accommodates full range of optics

DAC
- Designed for short distances, typically less than 5 m
- Low power consumption
- Low cost

CFP8
- Electrical interface supports both 400GAUI-8 and 400GAUI-16
- Form factor is approximately the same size as CFP2

EXFO solutions

Test in the lab, field or on the production floor

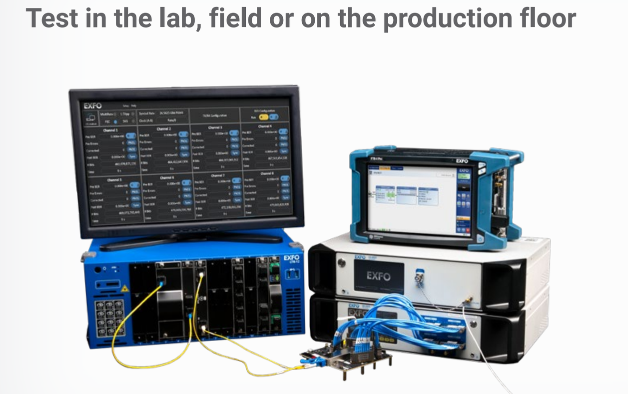

Open Transceiver System (OTS):
supports today’s and tomorrow’s optical interfaces

CFP8 QSFP-DD OSFP ....and ready for your next evolution.

EXFO