OIF
PLL INTEROP DEMO
OFC 2025

Energy Efficient Interfaces

OIF Interop Demos for EEI

Energy Efficient Interfaces (EEI) @ OFC 2025

- Energy Efficient Interfaces (EEI)
- EEI Interoperability Agreements
- Interoperability Demonstrations

What is needed?

AI compute connectivity is on an unsustainable path as it grows to larger clusters

https://www.energy.gov/ne/articles/5-facts-know-about-three-mile-island

What is needed:

- Copper for shorter reaches
- Optics for increased connectivity over longer distances
- Links that are:
  - Energy efficient links
  - Low latency links
  - High reliability links
- And these links must support a variety of protocols (Ethernet, PCIe, UEC, UALINK, NVLINK, IB, etc.)

Next generation links for AI compute

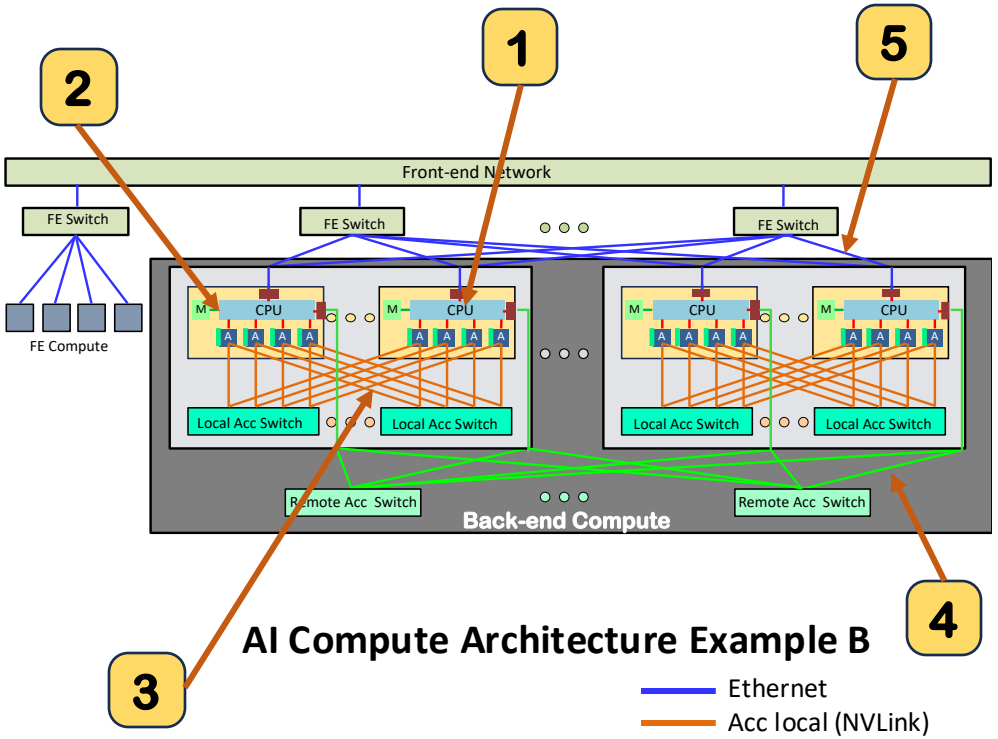

AI Compute Architecture Example B

- Ethernet
- Acc local (NVLink)
- Acc remote (InfiniBand)
- Memory I/F (CXL/PCIe)
- HBM I/F (Very wide I/F)
- Local connection (D2D)

Today, copper links are commonly used for scale-up links

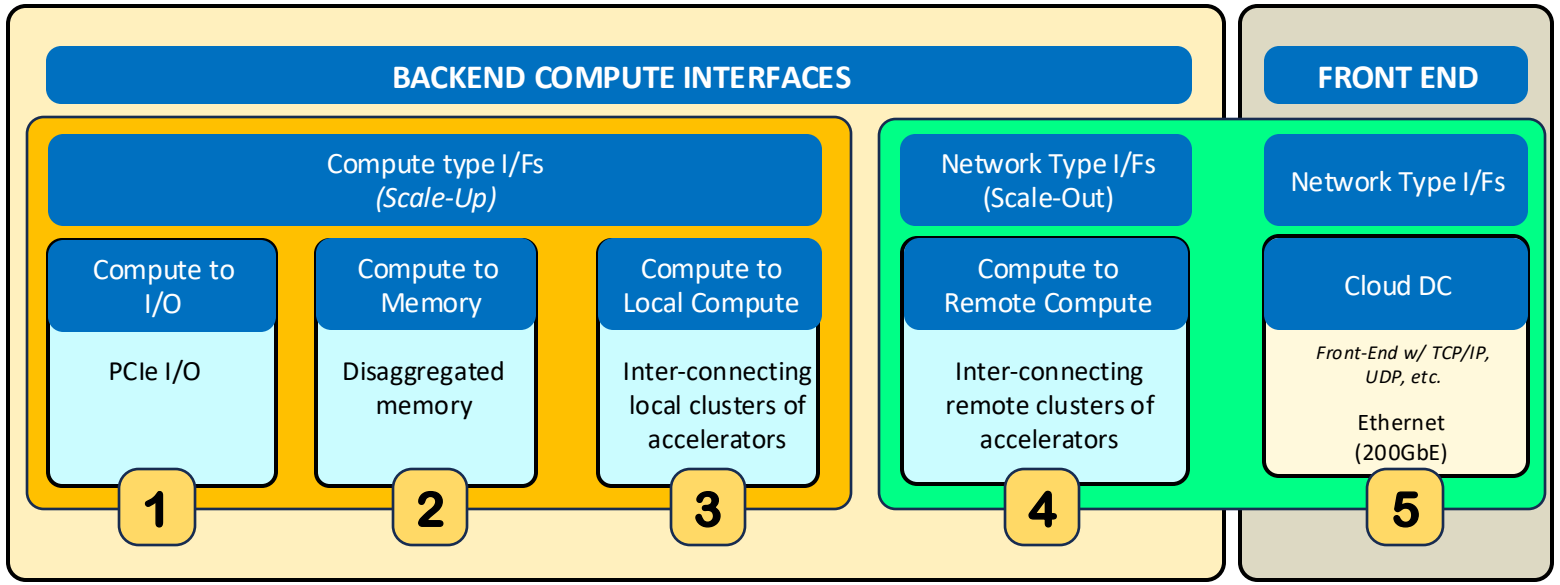

Scale-up "PCIe Like"			Scale-out & Front End "Ethernet Like"	
CRC with retry + very low latency FEC			***Ethernet FEC***	
PCIe	*(ex: 30x BW)* HBM	*(ex: 9x BW)* UPI NVLink UALink CXL.cache	*(ex: 1x BW)* InfiniBand Ethernet UEC	Ethernet
	COPPER & OPTICS CXL.mem		**OPTICS**	

Targets for next generation EEI links for AI compute

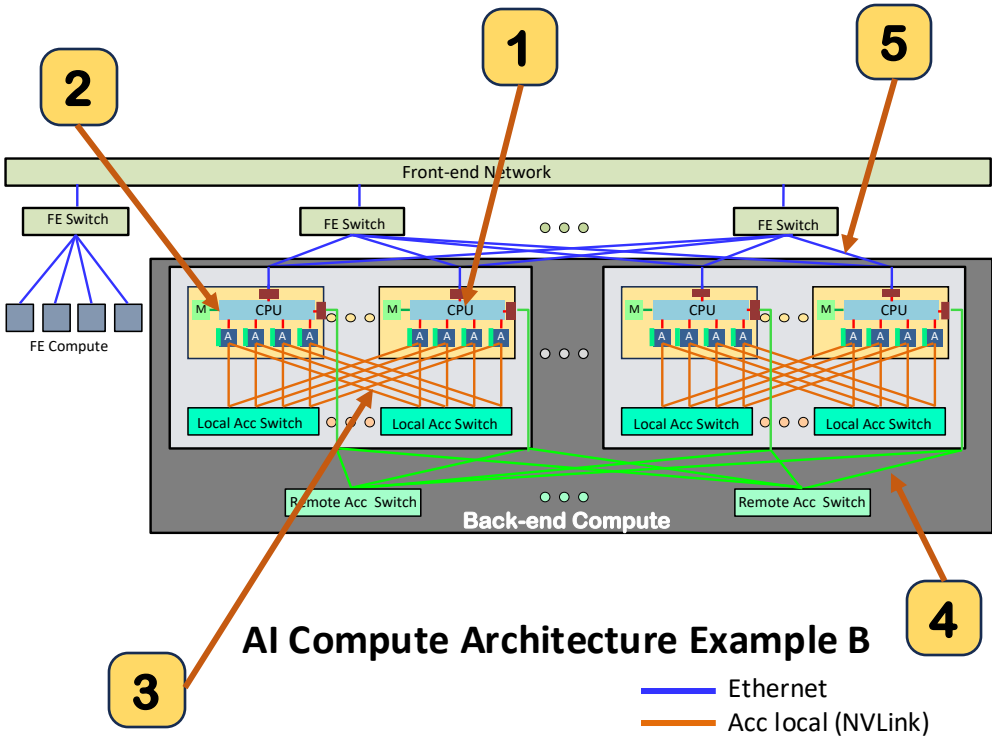

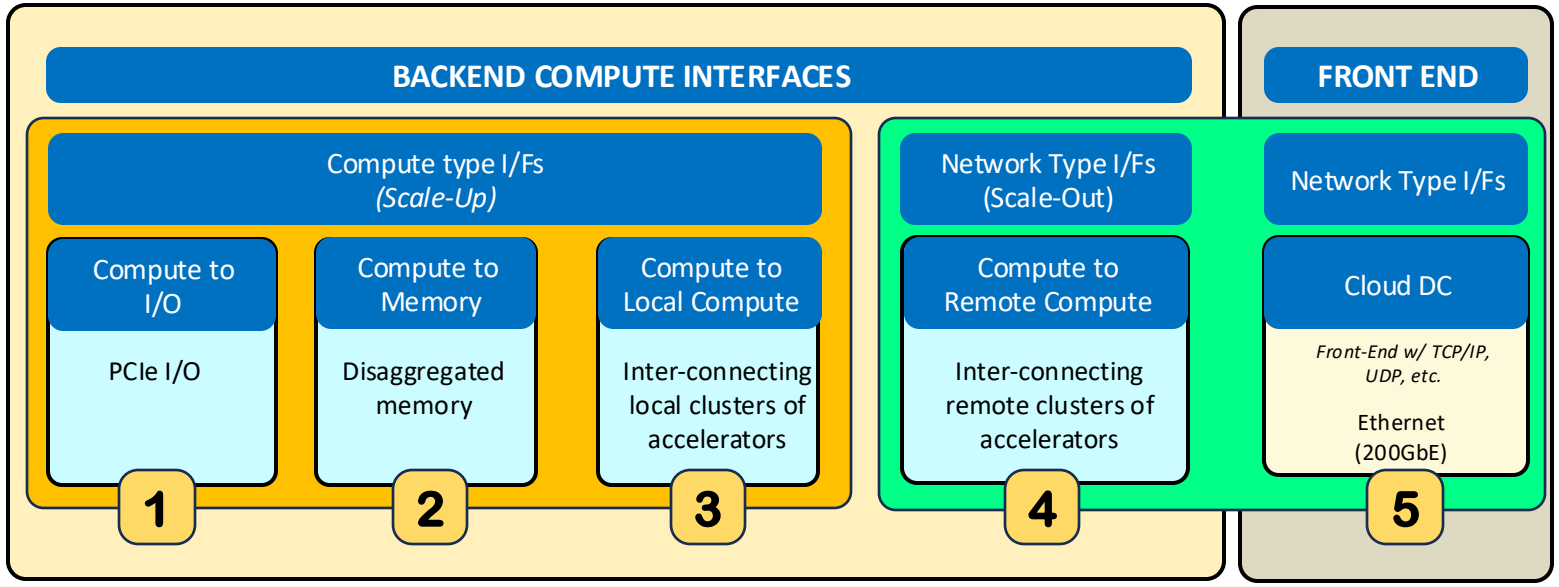

Parameter	Compute - Local	Compute - Network	Front - End
Scale	Local	10's or racks	Data center
Reach [m]	~ 10m	~100m	1km
BW Density [Tbs/mm] (Tx + Rx)	2.0 to 4.0	2.0 to 4.0	Std Ethernet
Latency[1] [ns]	< 5.0 + tof	< 5.0 + tof	< 20.0 + tof
Energy Utilization[1] [pJ/b]	< 5.0	< 10.0	< 10.0
Reliability (link errors)	high (low latency FEC w/ CRC)	high (lower latency FEC)	Std Ethernet
Reliability (hw failures)	high	high	Std

Hyperscaler members provide requirements for the next generation of EEI links

Approaches for EEI link interfaces for scale-up

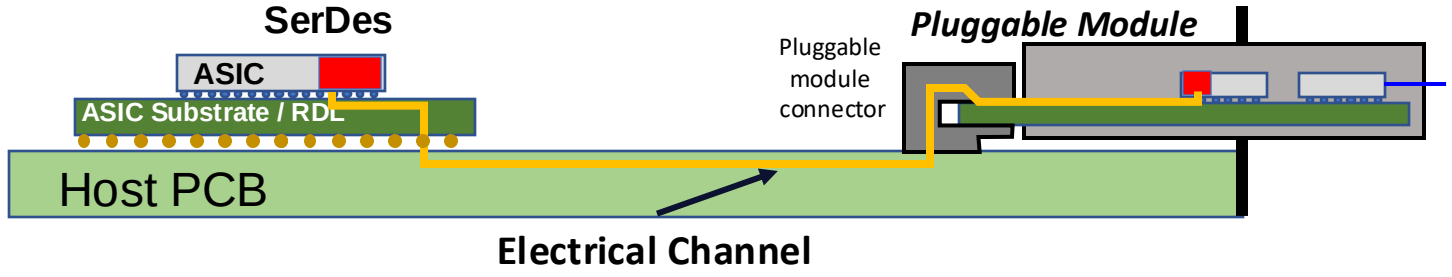

Pluggable

Replaceable modules:

- Non-retimed (e.g. LPO, LINEAR, passive copper)
- Tx-retimed (e.g. RTLR, LRO) (Tx diagnostics)
- Fully-retimed (full diagnostics)

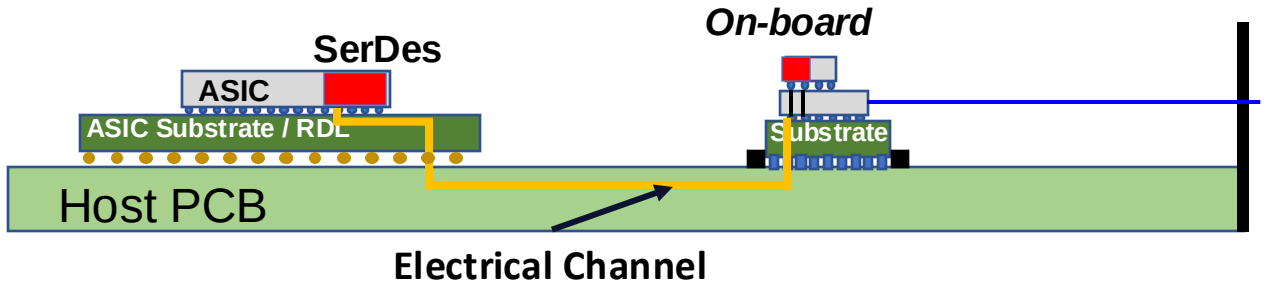

On-Board

- Increased density
- Shorter electrical channel
- Replacement requires card removal

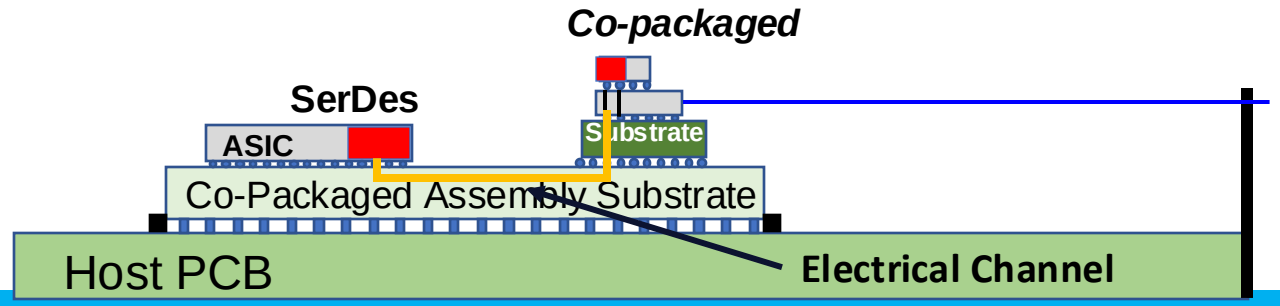

Co-packaged

- Highest density
- Shortest electrical channel
- Replacement requires repair of ASIC's package
- Configured during ASIC packaging

Energy Efficient Interface Activities at the OIF

OIF's Co-Packaging Projects

✓ **Co-packaging Framework Project**

OIF-Co-Packaging-FD-01.0 – Co-Packaging Framework Document

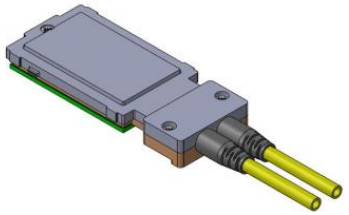

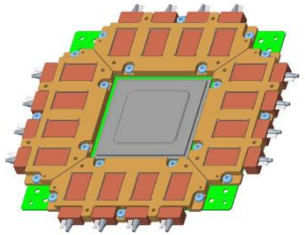

✓ **3.2T Co-packaged Engine**

OIF-Co-Packaging-3.2T-Module-01.0 – Implementation Agreement for a 3.2Tb/s Co-Packaged (CPO) Module

✓ **External Laser Source (ELSFP)**

External Laser Small Form Factor Pluggable (ELSFP) Implementation Agreement (August 2023)

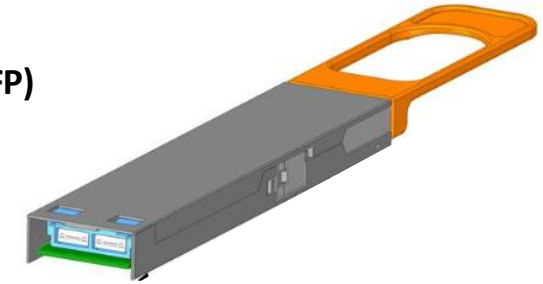

✓ **Management Interface for ELSFP**

OIF-ELSFP-CMIS-01.0 – Implementation Agreement for External Laser Small Form Factor Pluggable (ELSFP) CMIS

Energy Efficient Interfaces for AI

✓ **System Vendor Requirements for Energy Efficient Interfaces**

- Document the EEI requirements as provided by the end-users for AI/ML optical and electrical links

Energy Efficient Interface Framework

- Study and initiate new standards for dense, low power, low latency links for AI/ML

RTLR Project (Retimed Transmitter, Linear Receiver)

- Address lower latency and low power applications utilizing transmit retimed optical transceivers (e.g. Ethernet, UEC, etc.)

CEI-Linear (Non-Retimed Interface)

- Low power optical interface (LPO, CPO, & NPO)

COI Project (Compute Optics Interface)

- Address energy efficient, low latency photonic interfaces for transport of traffic for AI scale-up applications

High-Density Connector Project

- Defining requirements for next generation connectors

Summary

AI compute connectivity is on an unsustainable path as it grows to larger clusters

https://www.energy.gov/ne/articles/5-facts-know-about-three-mile-island

The OIF is working on the next generation of links:

- ✓ Copper for shorter reaches
- ✓ Optics for increased connectivity over longer distances
- ✓ Links that are:
  - Energy efficient links
  - Low latency links
  - High reliability links
- ✓ And these links must support a variety of protocols (Ethernet, PCIe, UEC, UALINK, NVLINK, IB, etc.)

Energy Efficient Interfaces (EEI) @ OFC 2025

- Energy Efficient Interfaces (EEI)

- EEI Interoperability Agreements
  - 3.2T Optical Module for Co-Packaging Project
  - ELSFP Project
  - Electrical Interfaces for EEI
    - XSR, XSR+
    - RTLR (Retimed Transmit Linear Receive)
    - Linear (Non-retimed) (a.k.a. LPO)

- Interoperability Demonstrations

3.2T Optical Module

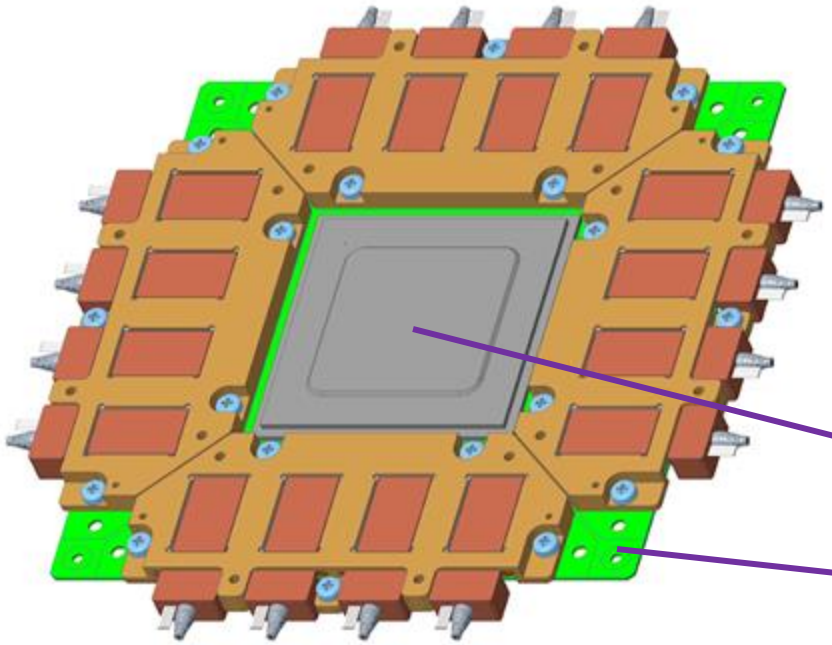

Co-Packaged Assembly Substrate (Interposer)

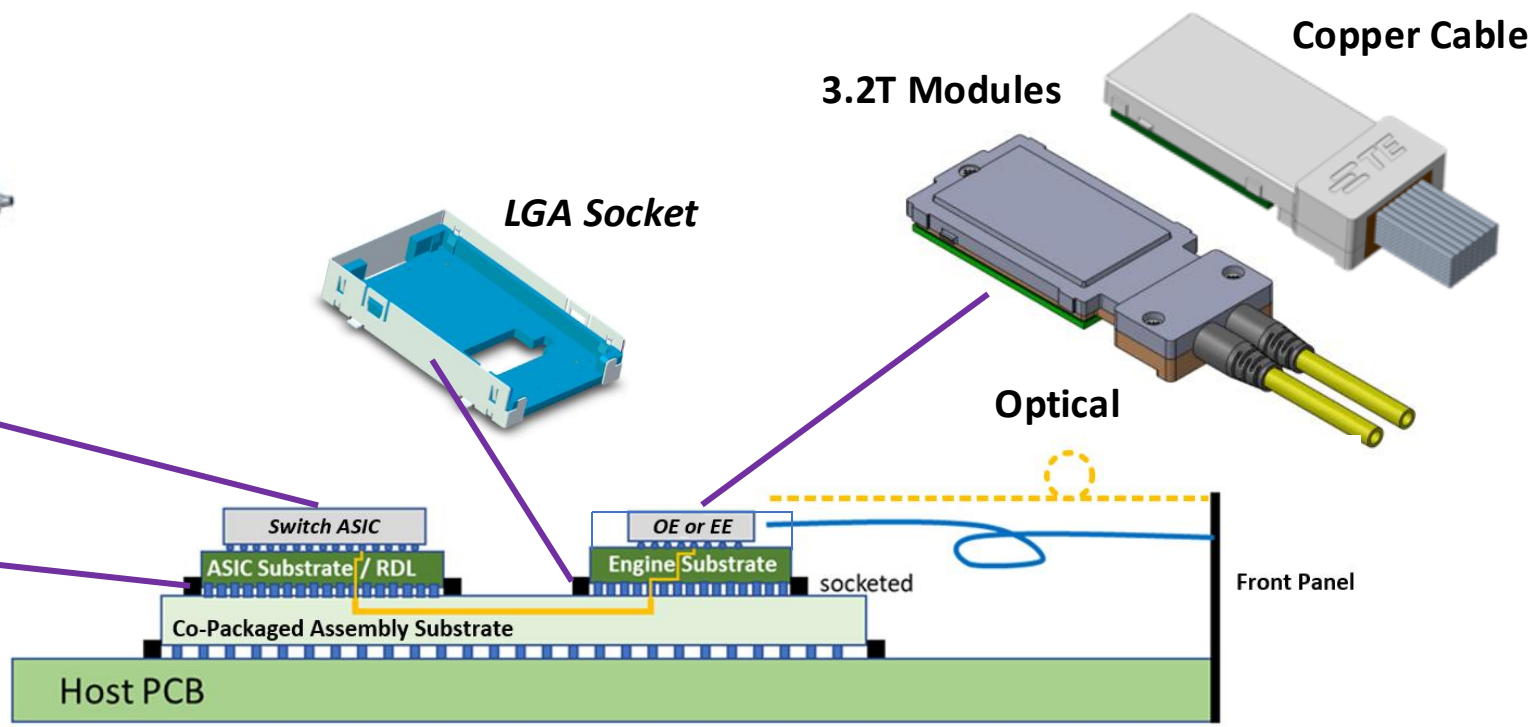

Channel components cross-section

3.2T Module Dimensions

- 32 x 112G XSR to Standard Optics:
  - 8 x 400G DR4
  - 8 x 400G FR4 (incl. 200G mode)
- Copper Cable Assembly compatible

- Power capability:
  - 56W (Internal Laser option)
  - 48W (External Laser option)

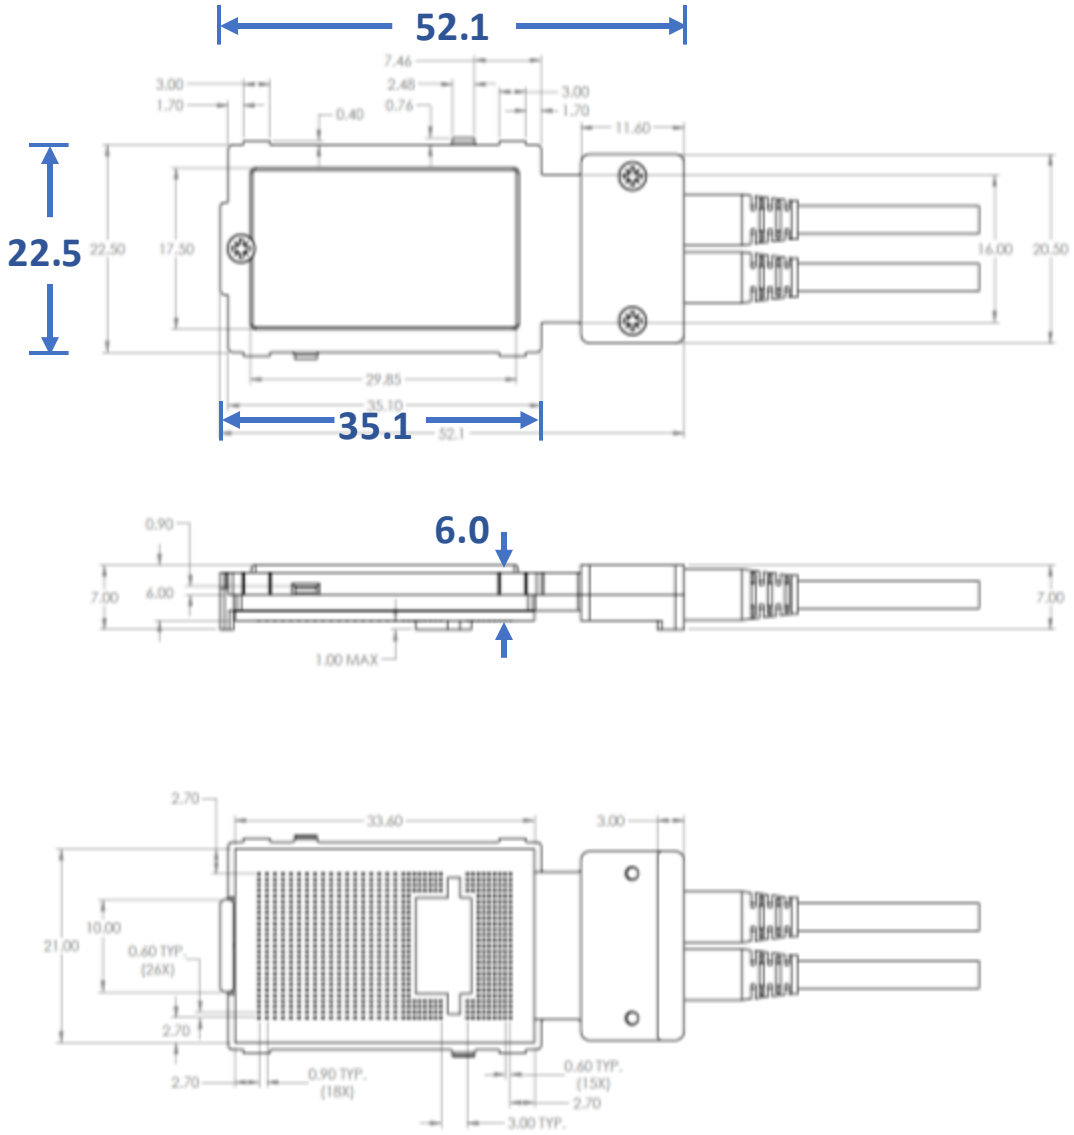

LGA Pin Map for 3.2T Module

- Supply rails:
  12V, 3.3V, 2.6V, 1.8V, 1.2V, 0.9V, 0.7V
- Comms Electrical: 1.2V SPI
- Comms protocol: CMIS
- 400G and 800G (2x400G) port grouping defined
  - *For low power modes and 2x400G-FR4 cable assignment*

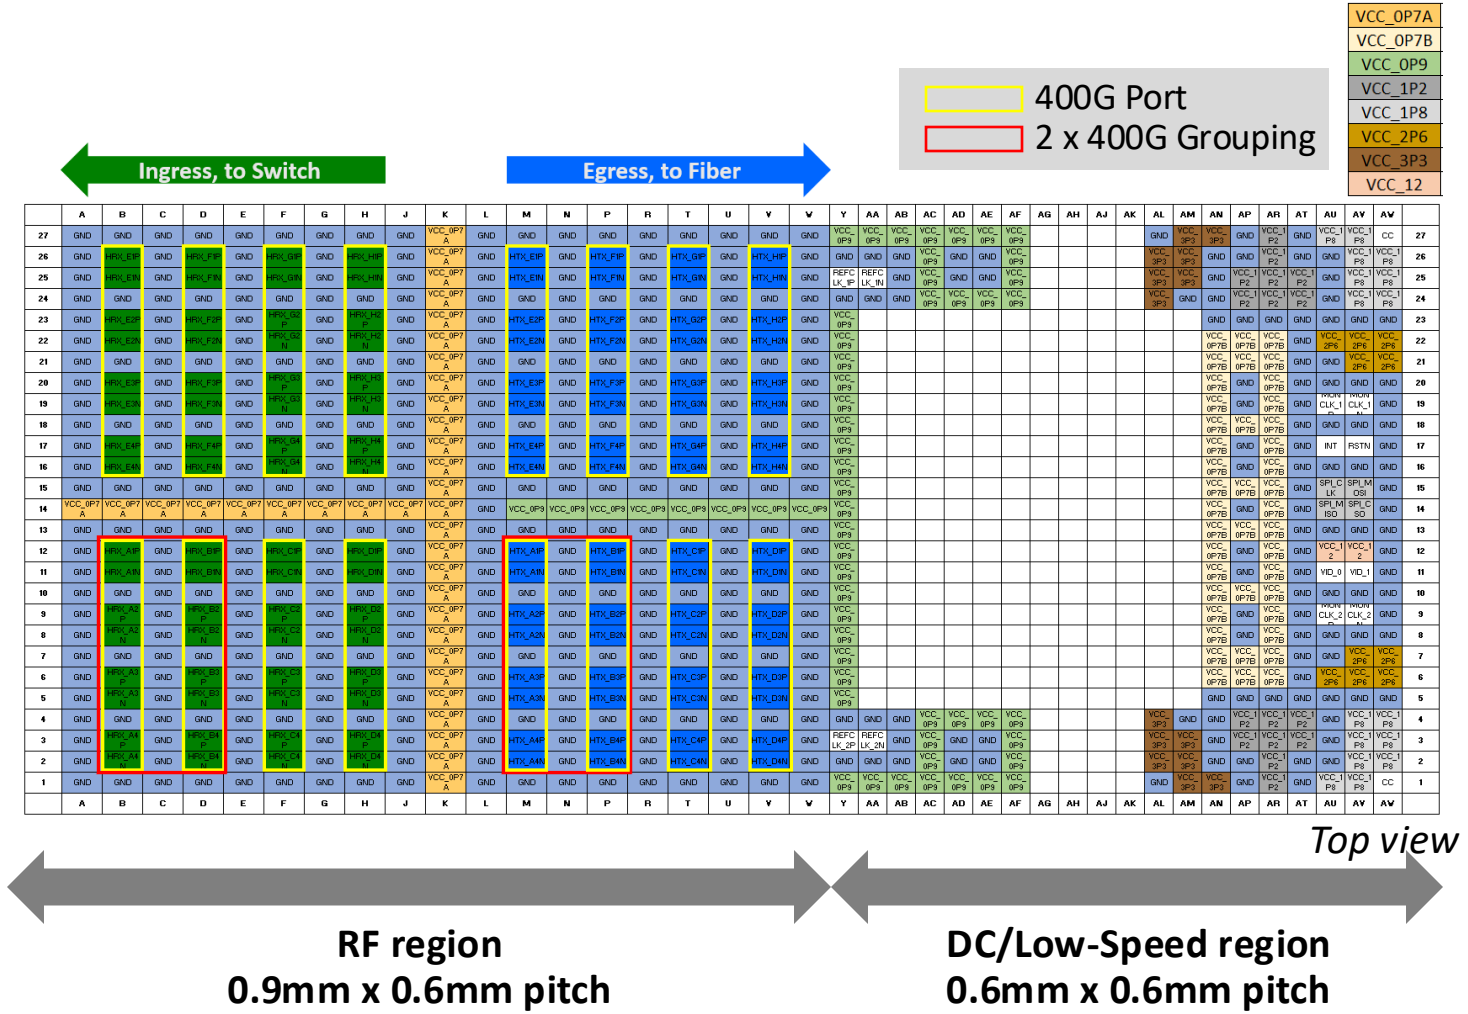

Energy Efficient Interfaces (EEI) @ OFC 2025

- Energy Efficient Interfaces (EEI)

- EEI Interoperability Agreements
  - 3.2T Optical Module for Co-Packaging Project
  - ELSFP Project
  - Electrical Interfaces for EEI
    - XSR, XSR+
    - RTLR (Retimed Transmit Linear Receive)
    - Linear (Non-retimed) (a.k.a. LPO)

- Interoperability Demonstrations

Why ELSFP?

- OIF defining common External Laser Pluggable
- Industry need for co-packaged and near-packaged systems
  - Systems need faceplate density
  - External laser modules need common specification for economies of scale
- Form factor to span multiple system generations
  - Plan for optical & thermal scaling

External Laser Small Form Factor Pluggable

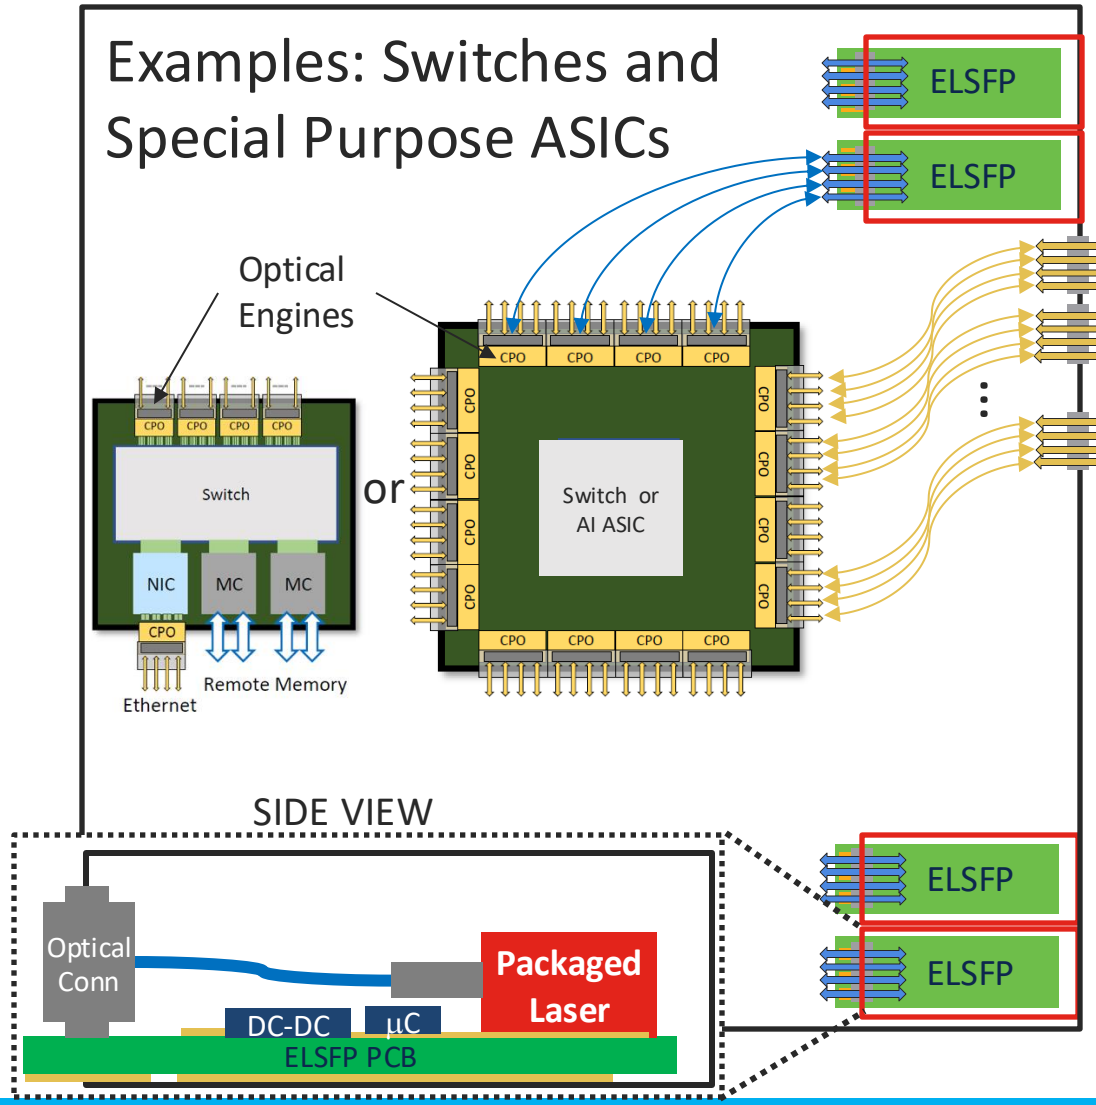

- ELSFPs provide CW laser power for optical engines (OEs).
- Decreases thermal power density in the system
- Each large system will likely need multiple (i.e. 8 or 16) ELSFPs
- The light from a given ELSFP can feed more than a single OE.
- A pluggable form factor helps to ensure total system reliability and a "hot swap" replacement if a single laser or ELSFP module fails.
- Eye safety is achieved by a blind mate optical connector internal to the system.

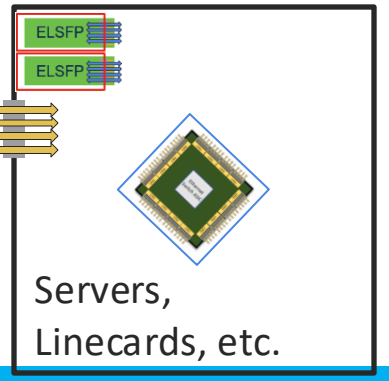

Initial Technical Concept

Density

- Blind mate pluggable
- Width similar to OSFP (16 modules wide with standard management I/O)

Commonality

- Industry standard 3.3V Supply
- CMIS (Common Management Interface Specification)

Scaling

- Optical power classes
- Thermal power classes
- Belly-to-belly configurations
- Riding heat sink for system flexibility
- Two “MT style” ferrules for future proofing
  - Support for 8 PM fibers per MT ferrule
  - One ELSFP to support multiple CPO modules

Single Port ELSFP Design

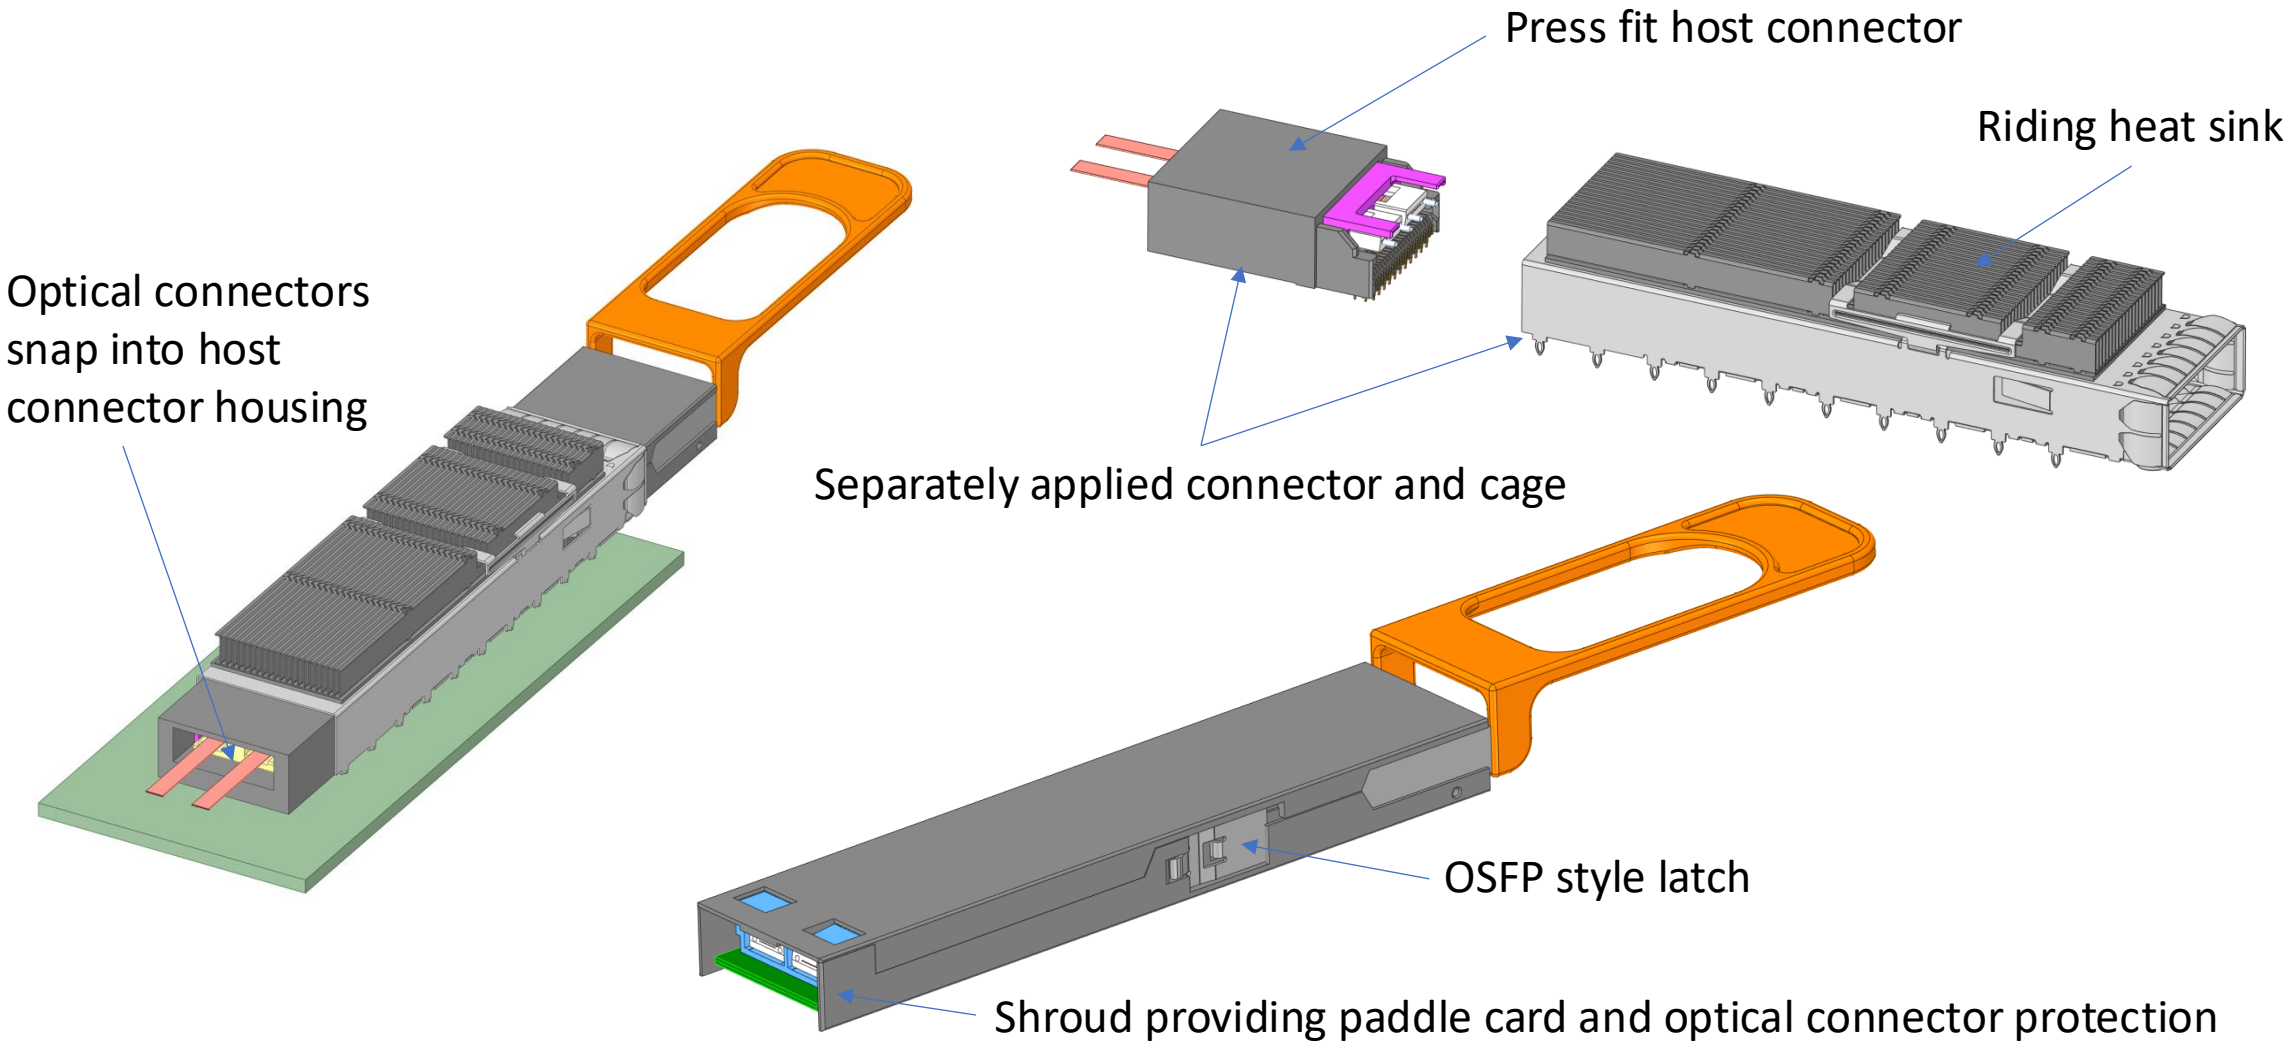

ELSFP Module-Side Optical Connector

Connector-to-Module Attachment (Optional)

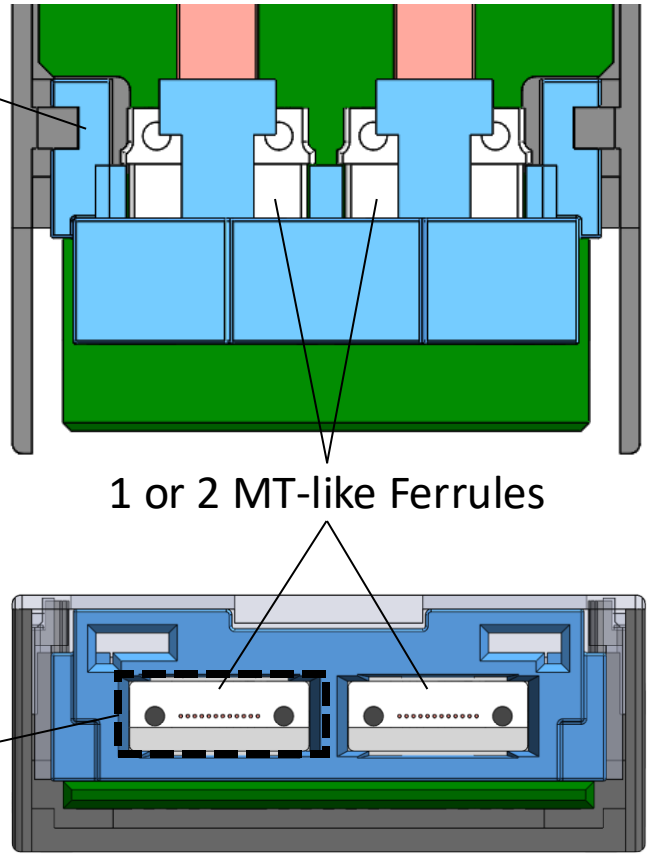

1 or 2 MT-like Ferrules

Optional 2nd MT-like Ferrule

Robust anchoring for optical connector

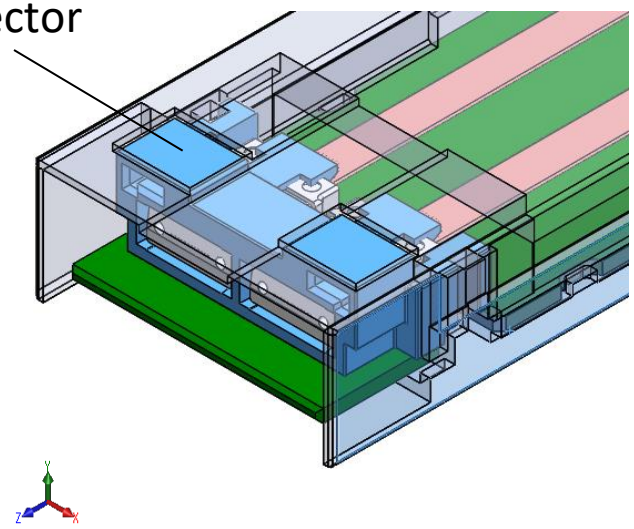

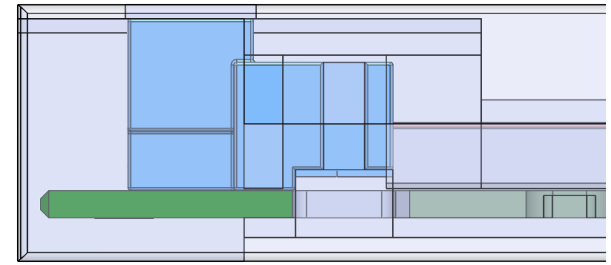

ELSFP Electro-Optical Connector

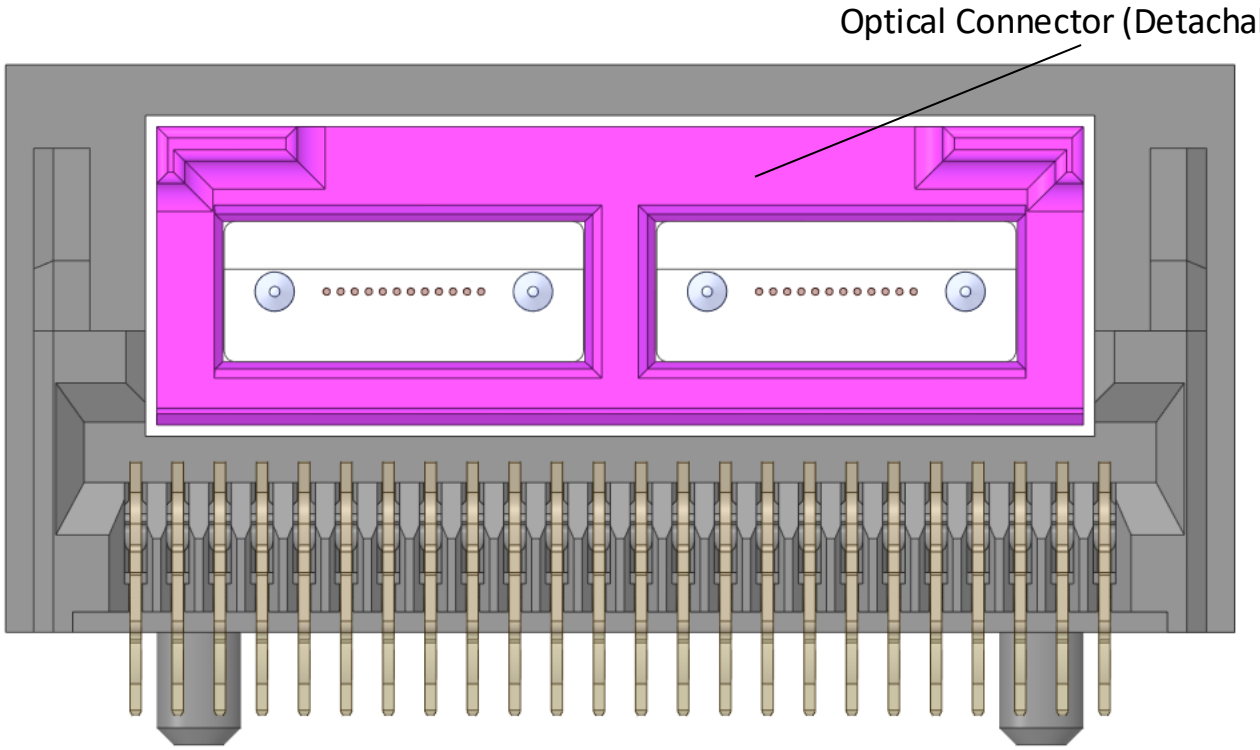

Host side Electro-Optical Connector

Module Bottom side Electrical Contacts

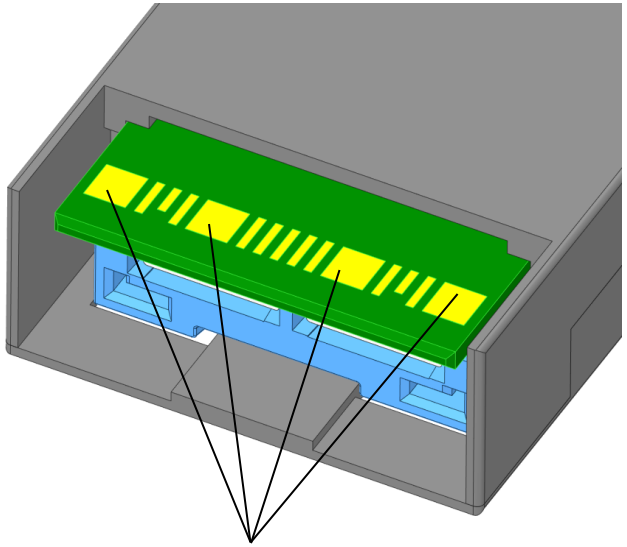

Sufficient power and ground to support **30W**

Additional pins for control/management, laser safety (i.e. presence pin), and spares for future proofing
Optical connector sub-assembly (pink) is separable from the board mounted electrical connector sub assembly

ELSFP Mating Sequence

STEP 1: Coarse alignment (PCB-to-host receptacle)

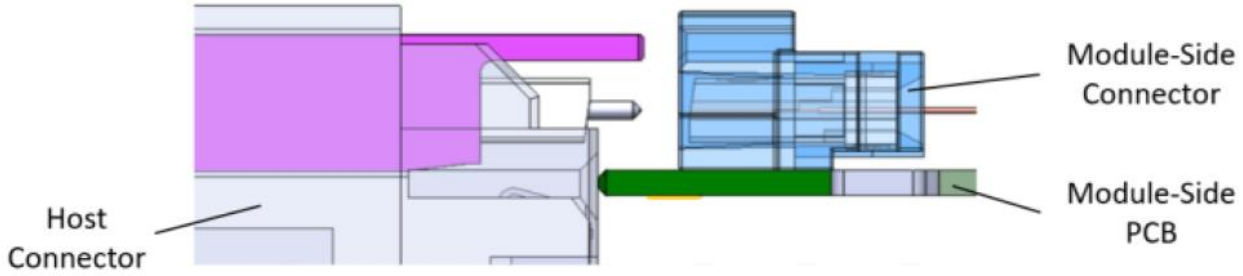

STEP 2: Coarse alignment (Optical coarse alignment pins)

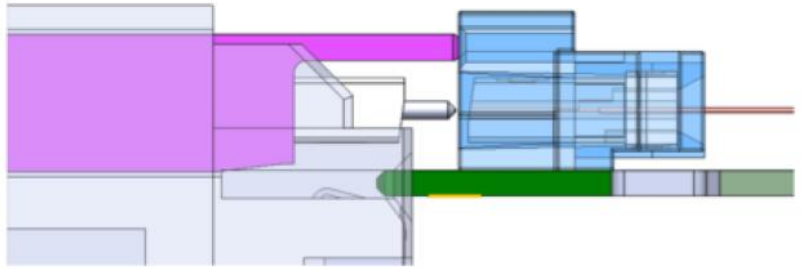

STEP 3: Fine alignment (ferrule guide pins)

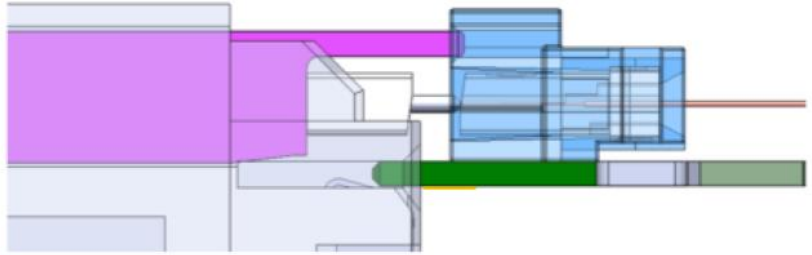

STEP 4: Ferrule end-faces in contact

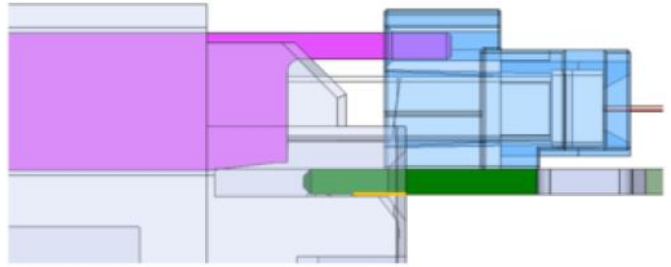

STEP 5: Electrical contact (presence pin)

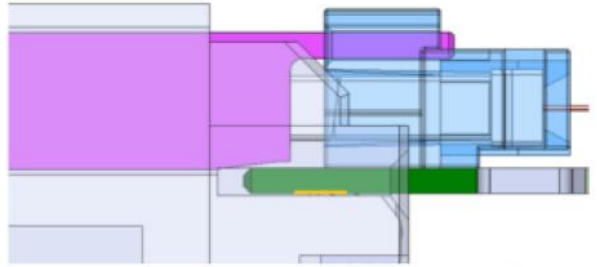

- Host side optical connector sub assembly has float to enable fine optical alignment.
- PCB and optical coarse alignment pins mate prior to fine alignment of optical ferrule guide pins .
- Ferrule end-faces to contact prior to electrical contact.

Optical Power Classes for ELSFP

ELSFP Optical Power Classes	Power/λ/Core +/- 1.5dB
Super Low Power - SLP	2dBm
Ultra Low Power - ULP	5dBm
Very Low Power - VLP	8dBm
Low Power - LP	11dBm
Medium Power - MP	14dBm
High Power - HP	17dBm
Very High Power - VHP	20dBm
Ultra High Power - UHP	23dBm
Super High Power - SHP	26dBm

Combs

Single-Channel

Multi-Channel

*Naming convention inspired by ITU Radio Frequency Band Nomenclature

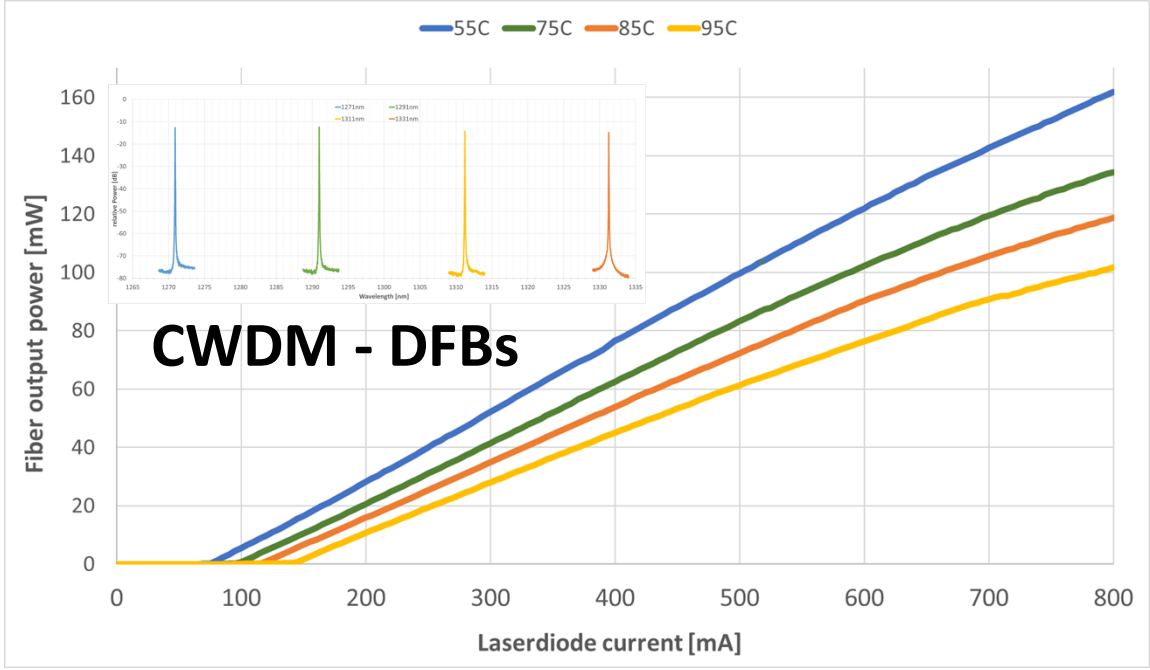

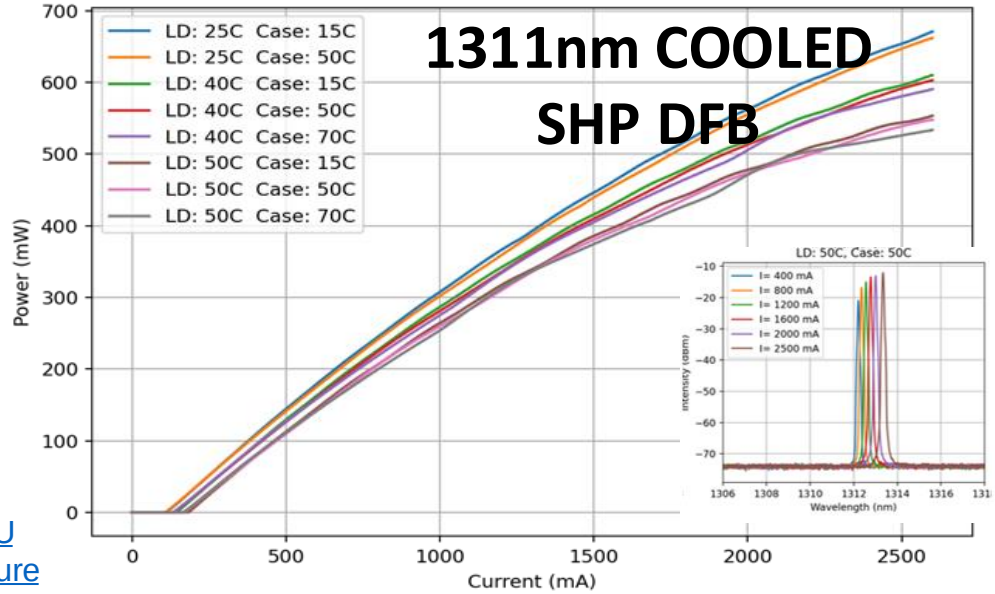

ELSFP's eco-system drives innovation

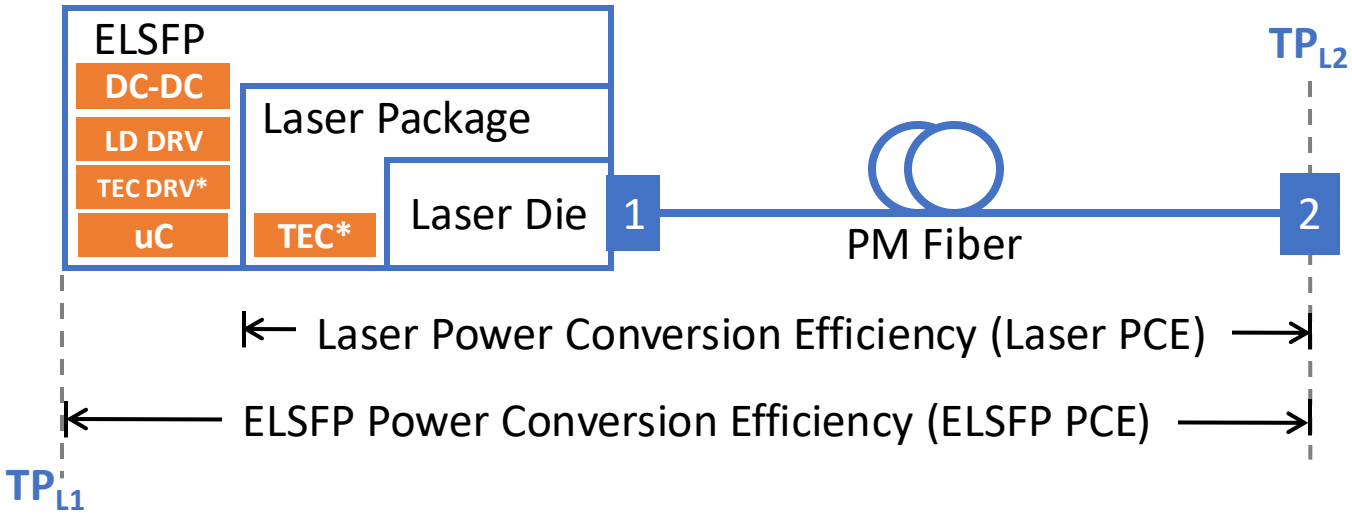

The ELSFP's eco-system continues to innovate and has yielded impressive improvements in energy efficiency (PCE), a key component of next generation energy efficient interfaces

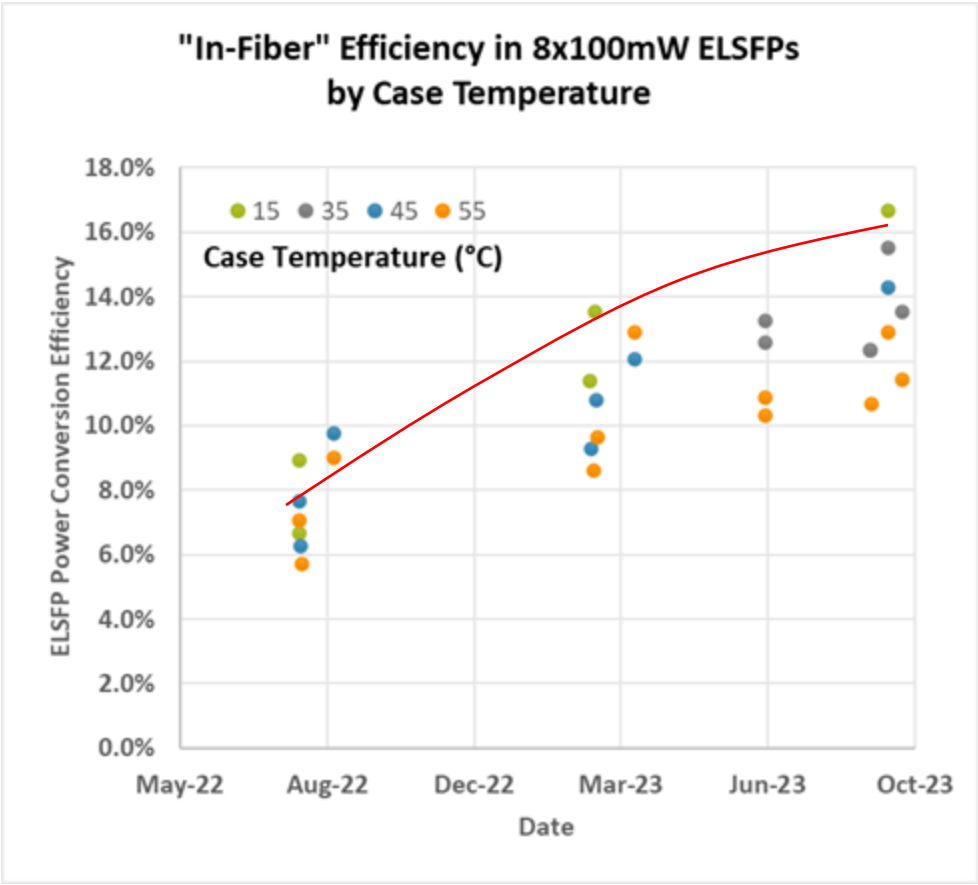

Eye Safety

ELSFP's blind mate optical connector paired with a system interlock enables a safer co-packaged system implementation for users.

Similar to EDFAs with powerful CW lasers, Class 3B and 4 lasers can be used inside ELSFP and systems can be deployed in unrestricted locations.

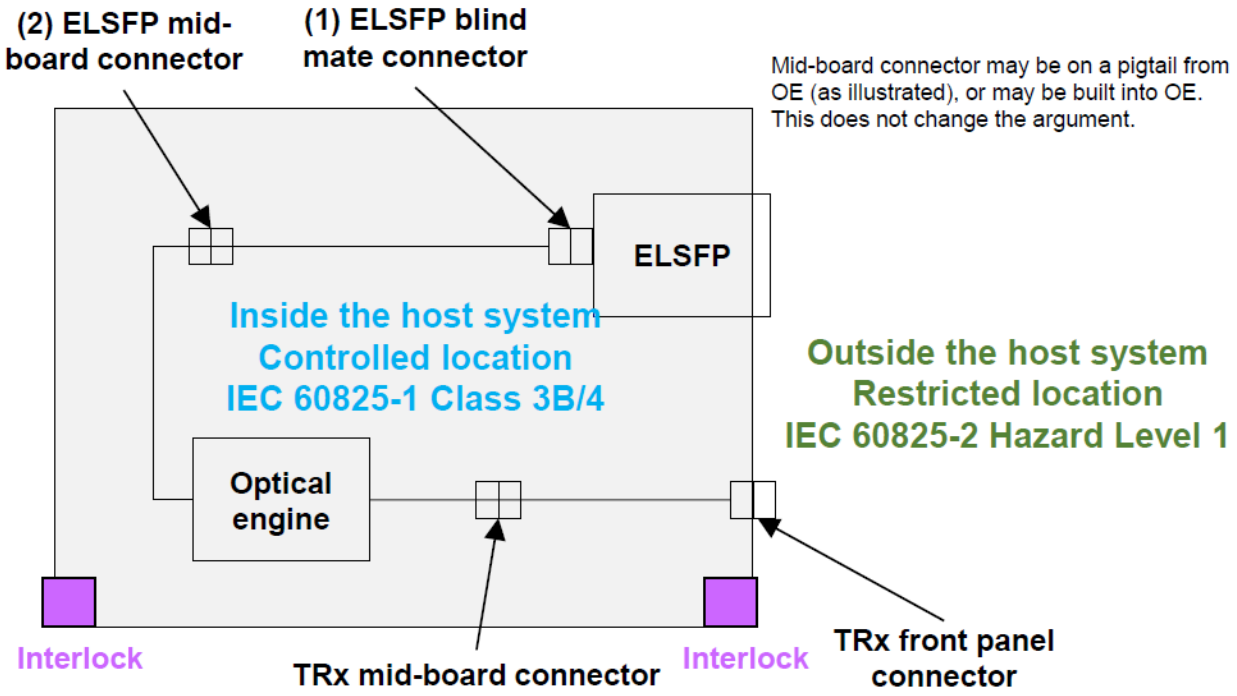

Energy Efficient Interfaces (EEI) @ OFC 2025

- Energy Efficient Interfaces (EEI)

- EEI Interoperability Agreements
  - 3.2T Optical Module for Co-Packaging Project
  - ELSFP Project
  - Electrical Interfaces for EEI
    - XSR, XSR+
    - RTLR (Retimed Transmit Linear Receive)
    - Linear (Non-retimed) (a.k.a. LPO)

- Interoperability Demonstrations

Electrical interfaces options for optical links

Retimed

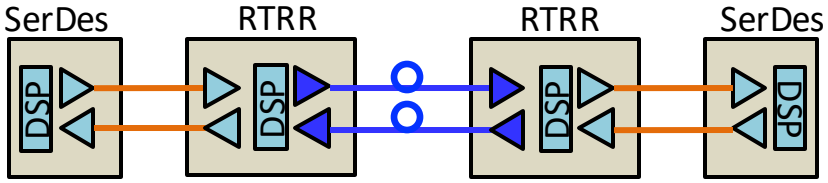

Fully Retimed Optical Link (Retimed Tx, Retimed Rx)

High power, Long electrical reach, Full diagnostics

Tx Retimed

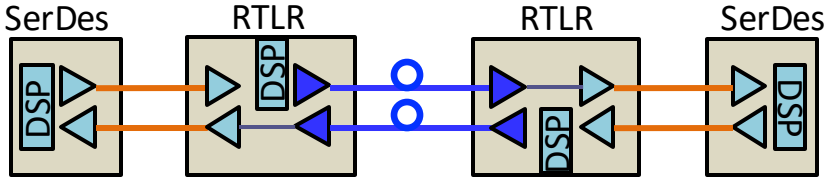

RTLR: Half-Retimed Optical Link (Retimed Tx, Linear Rx)

Balance of electrical reach with power, Full Tx diagnostics

Non-Retimed

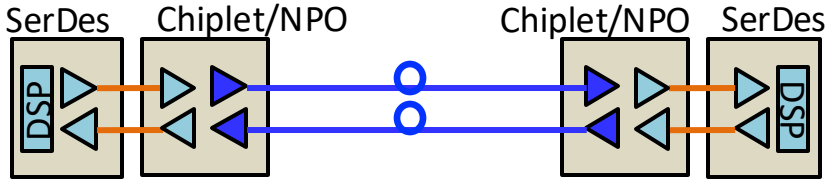

Linear: Non-Retimed Optical Link (Linear Tx, Linear Rx)

Low power, short electrical reach, Limited diagnostics

Slow Wide I/F

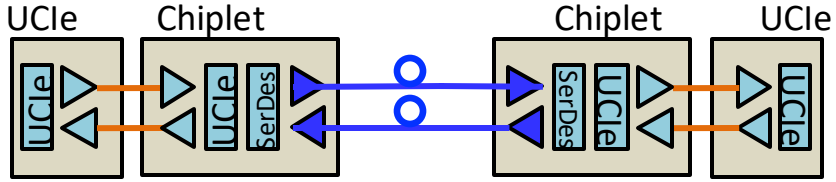

Co-packaged: Retimed Optical Chiplet Link Based Link

Low power, very short reach, Full diagnostics, not serviceable

CEI – An Essential Building Block for Co-packaging

Pluggable Module Channel Example Illustration

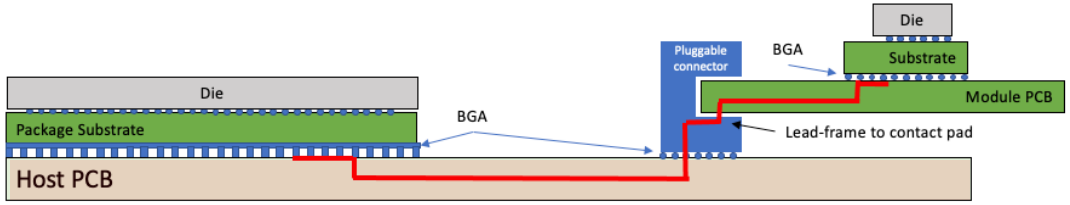

- Channel loss: 16dB ball to ball (22-24dB bump to bump)
- Typical pluggable connectors: IL of ~1dB with RL of -10dB @26.5GHz

CPO/NPO Channel Example Illustration

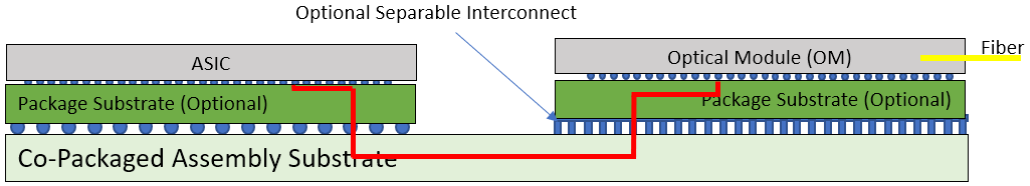

- Channel loss: CPO – 10dB bump to bump; NPO – 13dB bump to bump
- Optional separable interconnect performance example: LGA socket: IL of ~0.05dB with RL of -40dB @26.5GHz (*oif2020.341.01, Nathan Tracy*)
- Avoids/reduces major discontinuities.
- Optical modules are not end user pluggable.

- Significant power saving opportunity over VSR to be captured.
- A broad interoperable ecosystem is the key to success and can only be achieved through standardization.

CEI-112G-XSR-PAM4 for Co-packaging

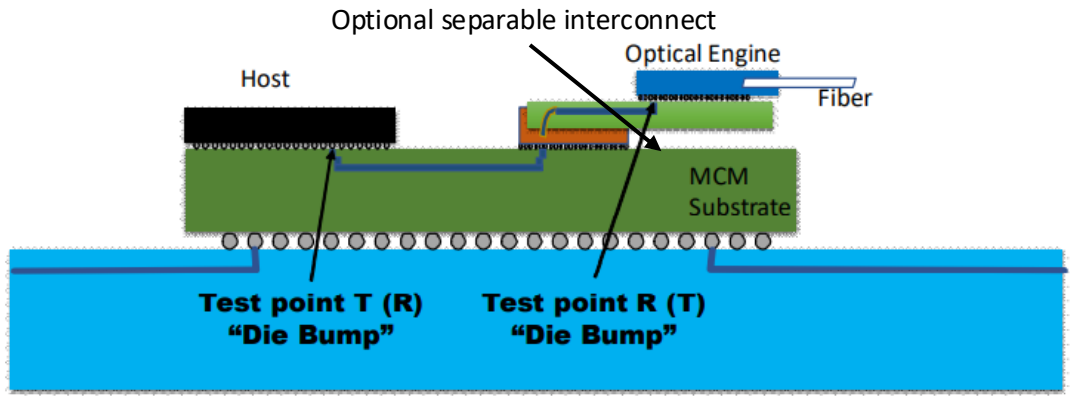

Category	IL at Nyquist (Max, dB)	BER (Max)
CAT1	10	1e-6
CAT2	10	1e-8
CAT3	8	1e-9

- Baud rates supported: 36 Gsyms/s to 58 Gsyms/s
- Based on loss and jitter budgets between TX and RX using copper signal traces in a SIP(System in a Package) to enable low power consumption
- Three channel categories are defined, allowing optimization for various applications.
- Timeline
  - Project started in April 2018.
  - Draft specification is becoming technically stable. Few pending items to be addressed.

CEI-112G-XSR+ -PAM4 for Near Packaging

- The emergence of Near Package Optics (NPO) Architecture
  - Co-packaging requires significant package substrate size increase and technology advancement, which adds risk to goals of availability, cost and multi-vendor support.
  - Instead of a monolithic package approach, Near Packaging relies on advanced PCB technology for dense high-speed routing without significant power penalty.
  - Near Packaging architecture takes advantage of existing technologies and more robustly enables an open ecosystem implementation.
- Additional margin also strengthens a broader supply base for co-packaging implementation and adoption.

- Baud rates supported: 36 Gsyms/s to 58 Gsyms/s
  - Optimize for Ethernet rate @ 106.25Gbps – the key application for CPO/NPO
  - Insertion loss < 13dB @ 26.5625GHz Nyquist bump to bump with up to 1 separable interconnect.
- Enable the lowest practical energy consumption (pJ/b) implementation.
- Leverage specification methodology and other work from existing CEI 112 projects.

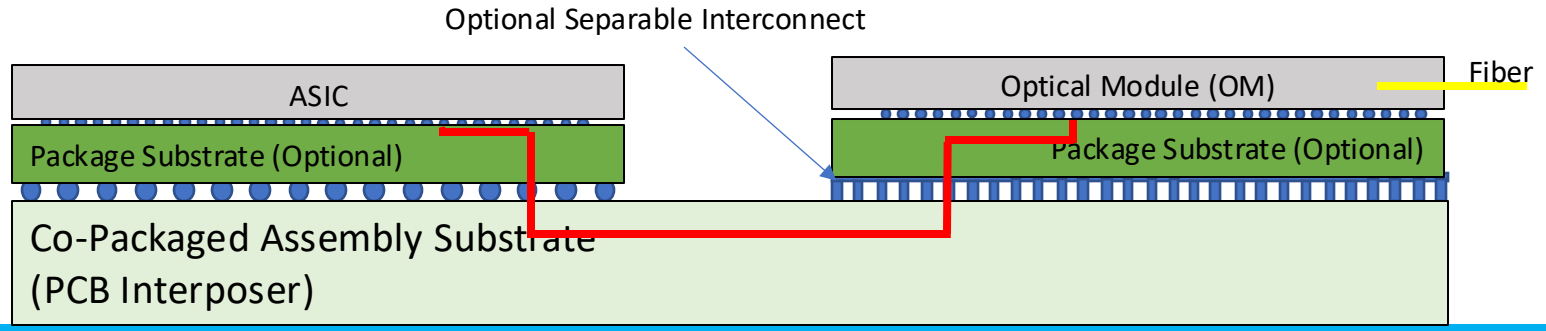

Optical Transceiver Retimer Permutations

Fully Retimed Optical Link *"DSP"*: Highest Power, Longest Reach

DSP/retimer module **OIF-CEI-112G-VSR-PAM4** supports 16 dB channel on egress with some optical output compliance expectation

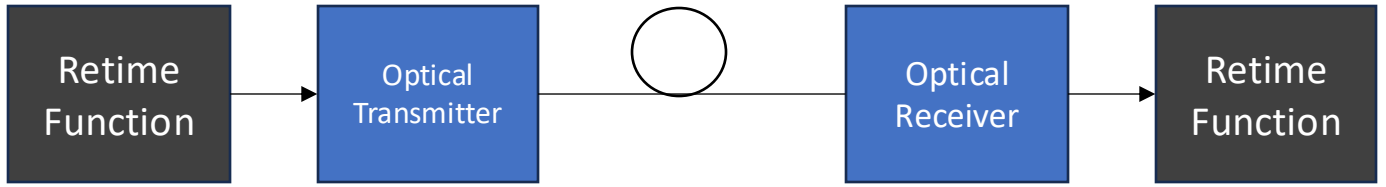

Ingress path includes DSP/retimer in the module and supports 16 dB channel to Host ASIC **(OIF-CEI-112G-VSR-PAM4)**

Retimed Transmit Linear Receiver (RTLR) *"LRO"* Optical Link: Balance of Reach, Power

OIF-EEI-112G-RTLR is meant to be identical to above on the egress channel **(OIF-CEI-112G-VSR-PAM4)**

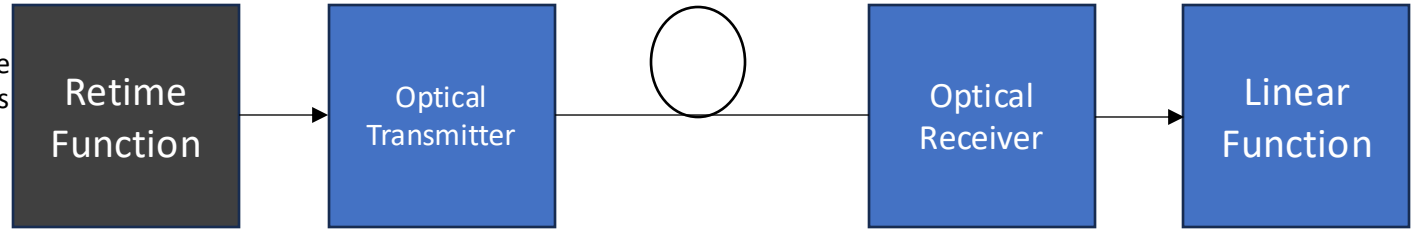

Ingress path removes the DSP/retimer in the module and uses an enhanced version of OIF **CEI-112G-Linear-PAM4** specifications by utilizing host ASIC DSP SerDes capability

Linear Non-retimed Optical Link *"LPO"* : Lowest Power, Shortest Reach

Egress path removes the DSP/retimer in the module and uses OIF **CEI-112G-Linear-PAM4** specifications by utilizing host ASIC DSP SerDes capability

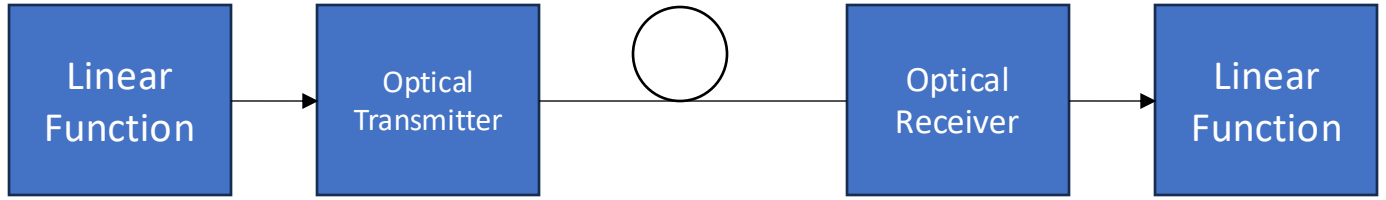

Ingress path removes the DSP/retimer in the module and uses OIF **CEI-112G-Linear-PAM4** specifications by utilizing host ASIC DSP SerDes capability

Energy Efficient Interfaces (EEI) @ OFC 2025

- Energy Efficient Interfaces (EEI)
- EEI Interoperability Agreements
- Interoperability Demonstrations

Interoperability Demonstrations:
External Laser Small Form Factor: ELSFP

Demonstrating 3 ELSFP modules showcasing the ecosystem

- Lasers: 8 lasers per module (1310nm)
- Output power: 20-25 dBm
- Both cooled and uncooled lasers

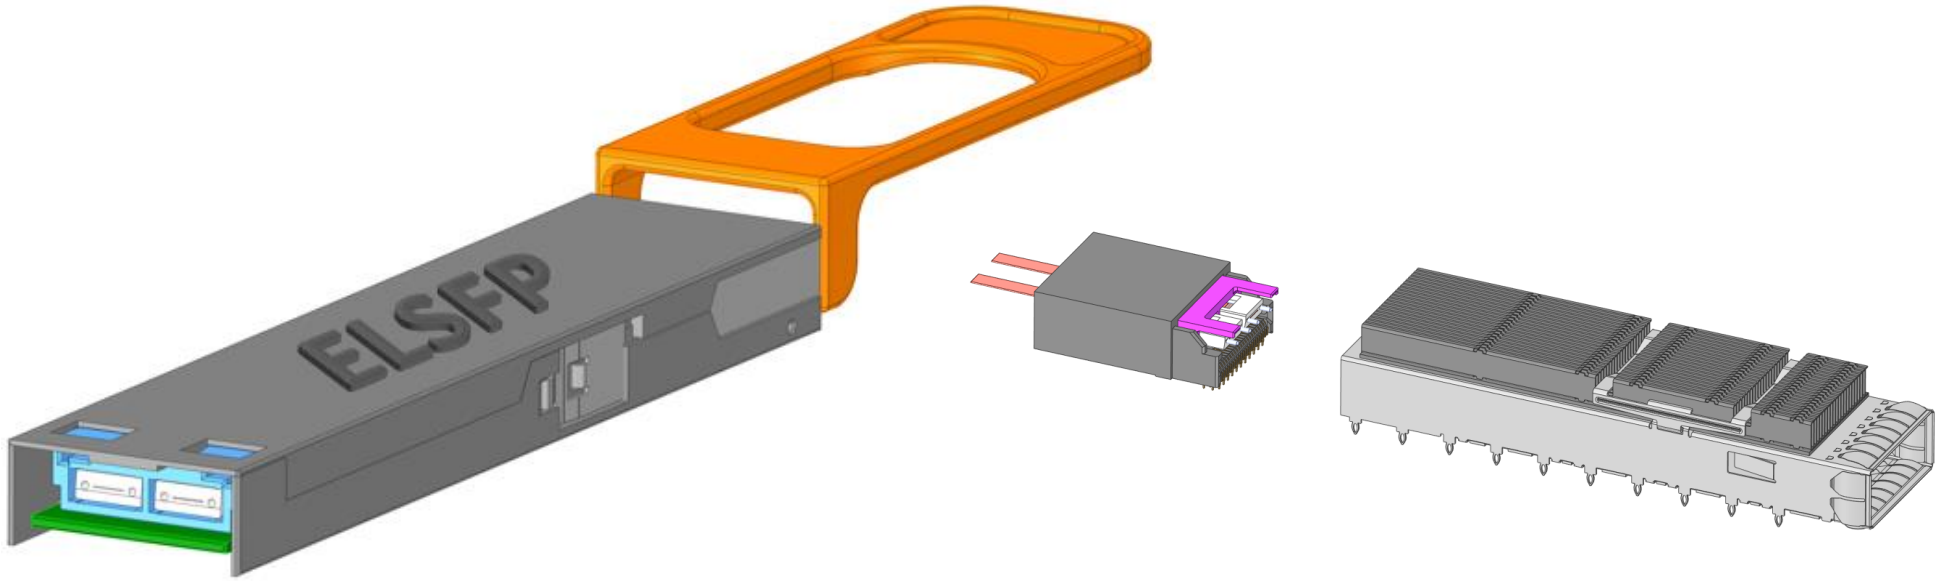

Interoperability Demonstrations:
External Laser Small Form Factor: ELSFP

Maintenance tools for optical connections

- Inspection tool
- Cleaning tool

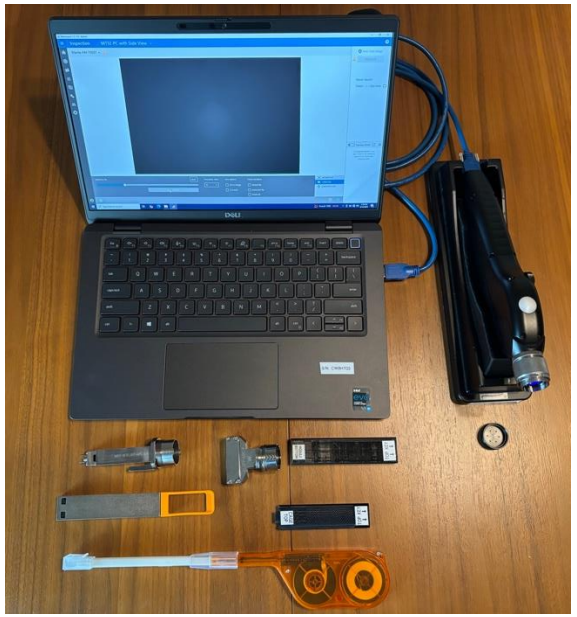

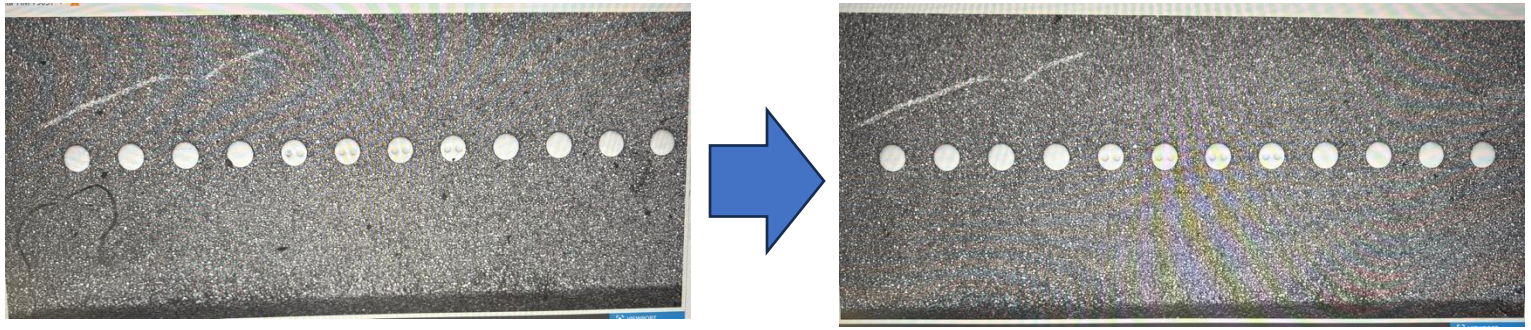

Before Cleaning

After Cleaning

Joint Demo: 112G RTLR & Linear Interoperability

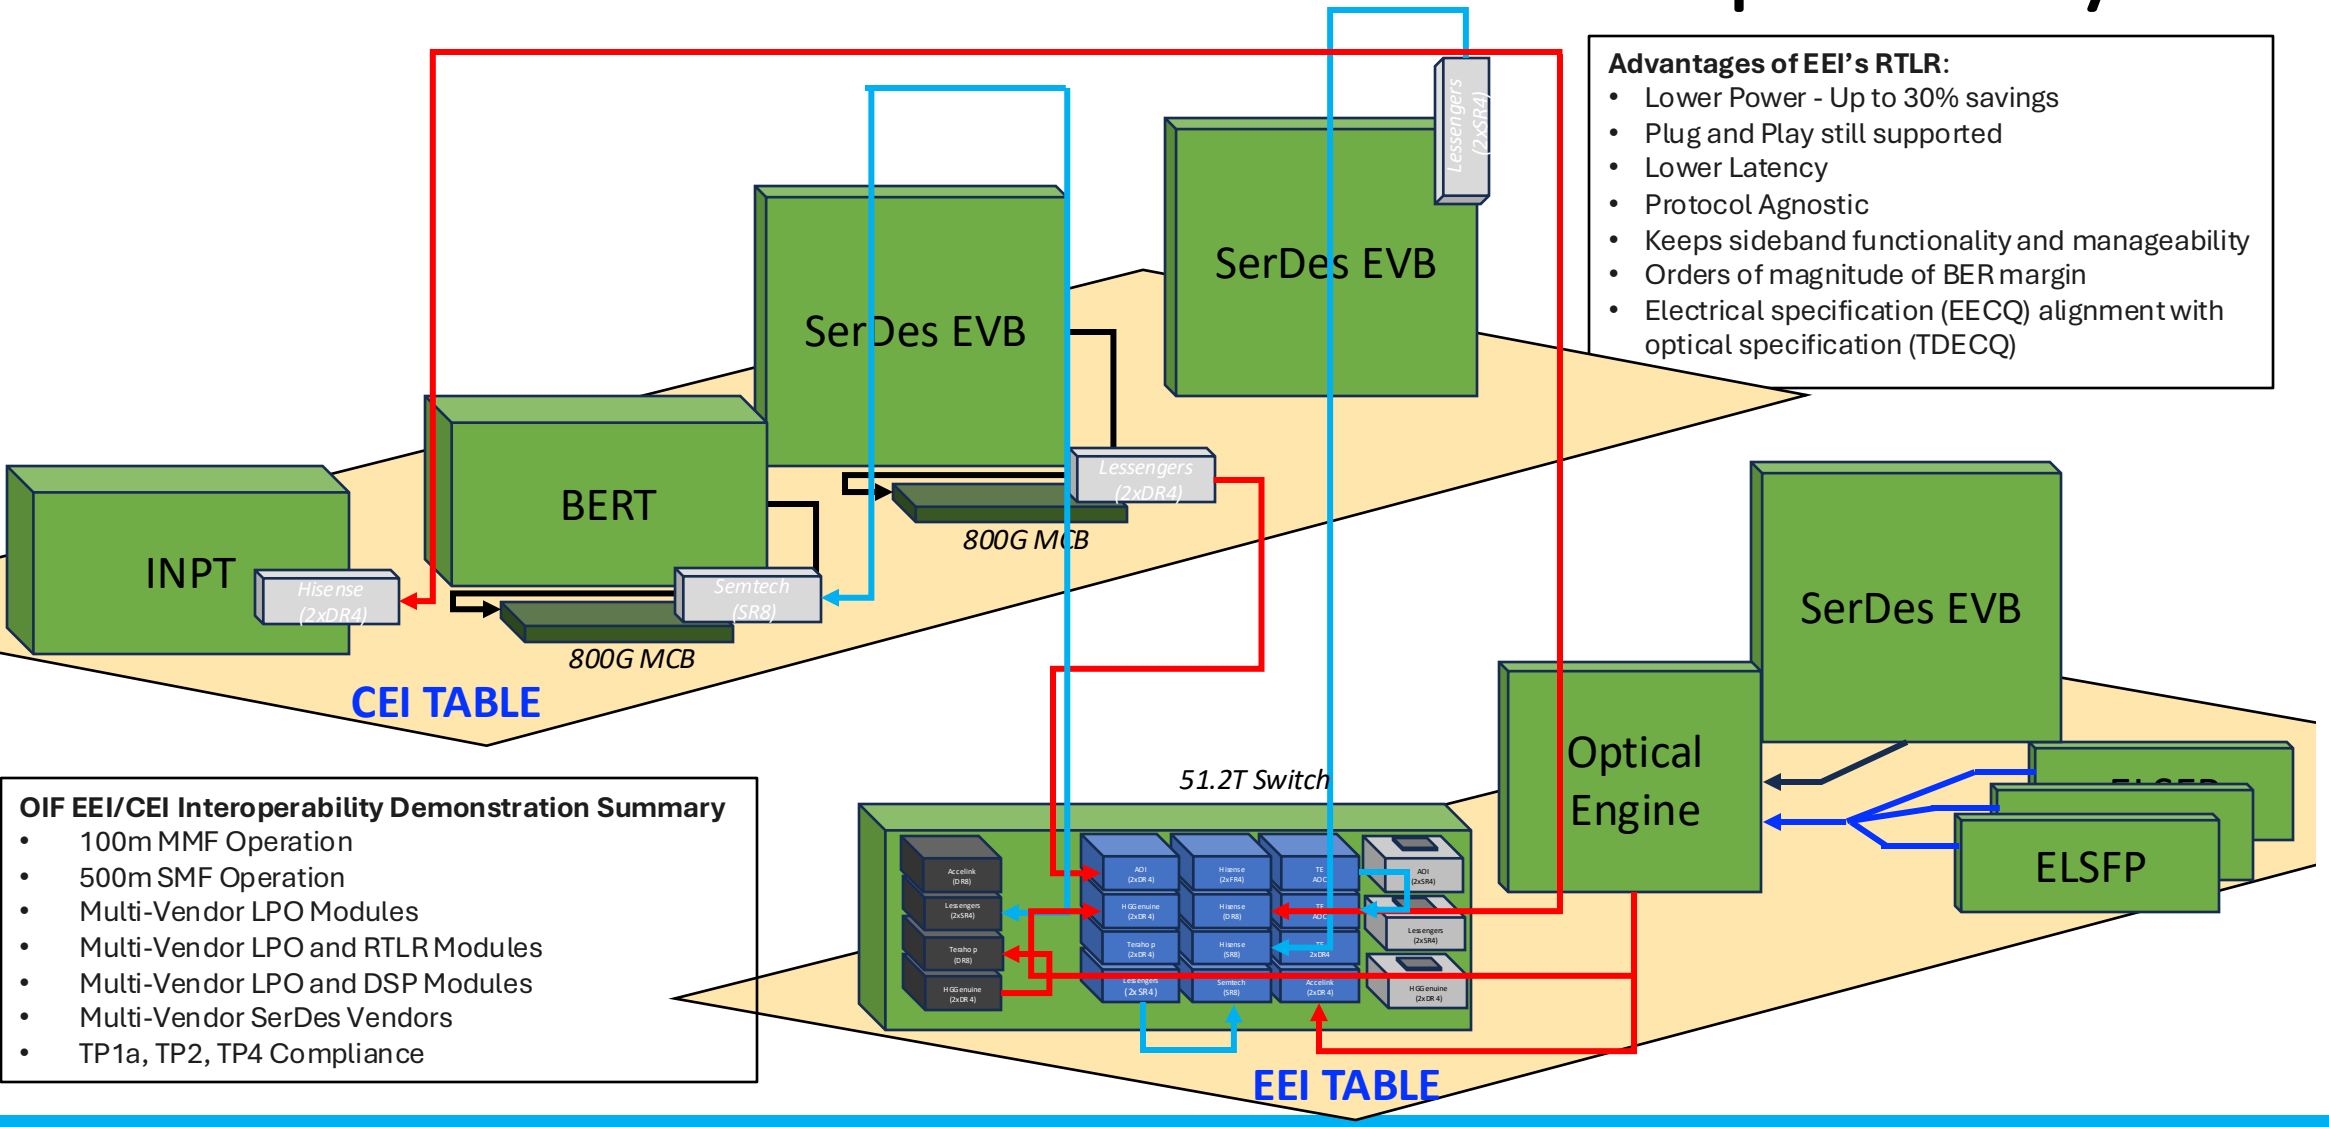

Advantages of EEI's RTLR:

- Lower Power - Up to 30% savings
- Plug and Play still supported
- Lower Latency
- Protocol Agnostic
- Keeps sideband functionality and manageability
- Orders of magnitude of BER margin
- Electrical specification (EECQ) alignment with optical specification (TDECQ)

OIF EEI/CEI Interoperability Demonstration Summary

- 100m MMF Operation
- 500m SMF Operation
- Multi-Vendor LPO Modules
- Multi-Vendor LPO and RTLR Modules
- Multi-Vendor LPO and DSP Modules
- Multi-Vendor SerDes Vendors
- TP1a, TP2, TP4 Compliance

ALPHAWAVE SEMI
CICT Accelink
Hisense Broadband
KEYSIGHT
LESSENGERS
MACOM
nubis COMMUNICATIONS
O-Net Technologies
SENKO Advanced Components
SUMITOMO ELECTRIC
SYNOPSYS
TE connectivity

EEI Conceptual Demo for AI Compute

Compute Chassis
Showing an array of compute and switch cards interconnected with a variety of optical connectivity options

AI backend compute employs low latency links to interconnect local accelerators in a cache coherent way. The local links are typically PCIe-like (NVLink, UALink, etc). Groups of compute clusters are interconnected with lower latency Ethernet / InfiniBand connections

Accelerator Cards
Variety of next gen PCIe compute cards plugged into a PCIe chassis

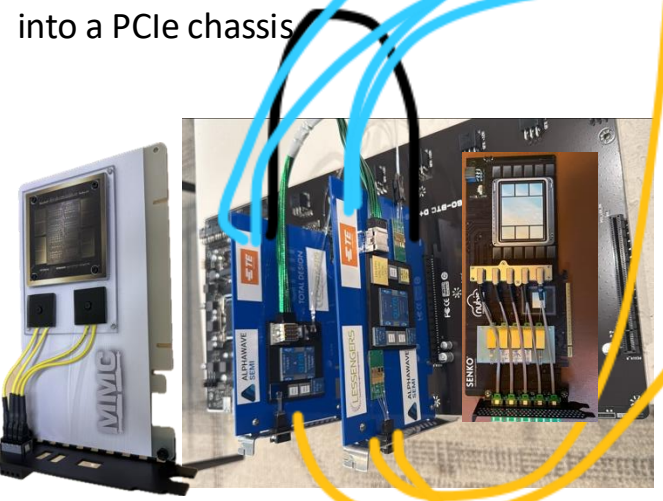

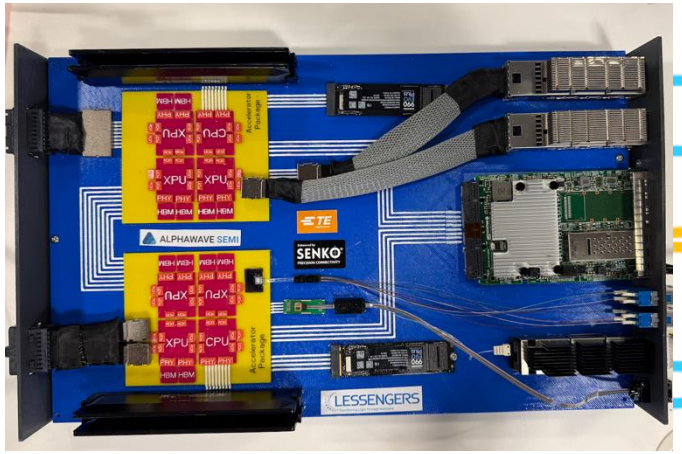

Next-Gen Switch Card located in compute chassis highlighting an ASIC with 4Tb/mm edge bandwidth and Ethernet interfaces on board. Optical links powered by ELSFP.

Accelerator Servers Next-gen AI compute blades heterogenous complex packages for dense scale out connections, along with integrated with fron access capabilities

OIF
ENERGY EFFICIENT INTERFACES

Thank-you