Next Generation CEI-448G Framework

OIF-FD-CEI-448G-01.0
November 4, 2025

Abstract:

As the OIF looks forward to the higher data rates and/or higher throughput required for the next generation of systems based on 448 Gbps per lane, a consensus has been reached that new specifications and technologies will be necessary. This framework document represents the OIF's efforts to identify hardware electrical interconnection application spaces where communication, compute, and AI systems could benefit from interconnection definitions or “Implementation Agreements” (IAs). The objective of this framework document is to identify key technical challenges for next-generation systems, define electrical interconnection applications, and discuss some of the interoperability test challenges. This will provide the OIF and other industry standards bodies with a common language and understanding of the development projects required for next-generation data rate systems.

Next Generation CEI-448G Framework Contributors:

Document Editors:

Cathy Liu	Broadcom Inc.
Nathan Tracy	TE Connectivity
Mike Li	Altera/Intel

Co-authors:

Cathy Liu	Broadcom Inc.
Nathan Tracy	TE Connectivity
Mike Li	Altera/Intel
Rob Stone	Meta
Kent Lusted	Synopsys
Yi Tang	Cisco
Jeffery Maki	Juniper
John Calvin	Keysight

Contributors:

Tad Hofmeister	Google
Ashwin Gumaste	Microsoft
Guangcan Mi	Huawei
Jim Weaver	Arista
Sam Kocsis	Amphenol
Bill Simms	Nvidia
Paul Brooks	Viavi

Table of Contents

List of Tables

List of Figures

Glossary†

ADC: An analog-to-digital converter is a system that converts an analog signal into a digital signal.

AI: Artificial intelligence.

ACC: Active copper cable. ACCs are a type of active copper cable used in data centers. ACCs primarily use Redriver chips and Continuous Time Linear Equalization (CTLE) to amplify and equalize the signal and provide longer reach than DAC's. ACCs generally are lower power and cost than AECs (another type of active copper cable), but provide shorter reach than AECs.

AEC: Active electrical cable. AECs are a type of active copper cable used in data centers. AECs utilize Retimer chips with Clock and Data Recovery (CDR) to reshape and retransmit signals which generally offer longer reach and better signal integrity compared to ACCs, but at a higher cost.

Application spaces: Portions of equipment or network architecture that could benefit from having a defined set of interconnection parameters.

ASIC: An application-specific integrated circuit is an integrated circuit (IC) customized for a particular use, rather than intended for general-purpose use.

Backplane: A group of electrical connections used as a backbone to connect several printed circuit boards together to make up a switch, computing or storage system.

BER: Bit error ratio is a measure of the number of bit errors that occur during data transmission, expressed as a ratio of erroneous bits to the total number of bits transmitted.

BoW: Bunch of wires is a die-to-die interface specification. It is part of the open compute project (OCP) and, like UCIe, is used to connect chiplets within a package.

Cabled host: An implementation where twinaxial cable is used instead of PCB traces expressly for the insertion loss benefits or 3D architecture.

CDR: Clock and data recovery, a component that re-establishes the timing of a signal that may have degraded due to impairments on a transmission line, the retimed signal is able to continue further to its destination.

CEI: Common Electrical IO, an OIF Implementation Agreement containing clauses defining electrical interface specifications, each optimized for various reaches at minimal power.

Coded modulation: Coded modulation is a technique in digital communication that combines error-control coding with modulation to enhance reliability and efficiency. Its main goal is to balance bandwidth efficiency, power efficiency, and error probability.

CMIS: Common management information specification.

CMIS-LT: CMIS-based link training provides a message set and exchange mechanism for out-of-band link training or tuning between a pluggable module and a host.

CMIS-VCS: CMIS versatile control set extends the base CMIS standard to allow for more advanced and flexible signal integrity (SI) capabilities.

CPC: Co-packaged copper is an emerging interconnect technology where copper cables are directly attached to the top of an Application-Specific Integrated Circuit (ASIC) or other high-speed integrated circuit within a single package. This design minimizes the length of high-speed electrical signals on the printed circuit board (PCB) and package, addressing the limitations of traditional copper traces at increasingly high data rates (e.g., 224G and above).

CPO: Co-packaged optics. An electrical to optical device intended to be mounted on the host package.

CTLE: Continuous time linear equalizer.

DAC: Direct attach copper cable. A high-speed cable assembly made of copper twinaxial cable with fixed passive transceiver modules on each end. DAC cables enable direct, electrical connections between networking devices over short distances.

DER: Detector error ratio.

DFE: Decision feedback equalizer. An equalizer by adding a filtered version of previous symbol estimates to the original filter output.

DSP: Digital signal processing.

ENOB: Effective number of bits. A measure of the dynamic performance of an analog-to-digital converter (ADC).

Faceplate: A plate, cover, or bezel on the front of a device which may contain I/O ports.

FEC: Forward error correction gives a receiver the ability to correct errors without needing a reverse channel to request retransmission of data.

FFE: Feed forward equalizer.

Gbps: Gigabits per second. The throughput or data rate of a port or piece of equipment. Gbps is $1x10^9$ bits per second.

GBd: The baud rate is the number of electrical transitions per second, also called symbol rate. Giga Baud is $1x10^9$ symbols per second.

Gearbox: A component used for managing and manipulating data streams, primarily by converting multiple serial data streams at one rate to multiple streams at another rate,

Hamming codes: A type of linear error-correcting code used in digital communication and data storage systems. It enhances data integrity by detecting and correcting single-bit errors that may occur during transmission or storage.

HPC: High performance compute.

HPDC: High-performance data center. A general term for data centers designed for high-compute workloads like AI/HPC.

IA: Implementation Agreements, what the OIF names their defined interface specifications.

IC: Integrated circuit.

IMDD: Intensity modulation direct detection is a method where the intensity of a light source is modulated by an electrical signal. This modulated light then travels through an optical medium (like a fiber optic cable) and is directly detected by a photo detector at the receiving end. This is a common and relatively straightforward technique for transmitting information over optical links.

I/O: Input Output, a common name for describing a port or ports on equipment.

ISI: Intersymbol interference.

KP4 FEC: A specific Reed-Solomon FEC (544,514) defined in IEEE 802.3 Clause 91, commonly used in Ethernet standards.

LDPC: A low-density parity-check code is a linear error correcting code, a method of transmitting a message over a noisy transmission channel. An LDPC is constructed using a sparse Tanner graph. LDPC codes are capacity-approaching codes.

LR: Long reach. CEI LR specifies backplane/midplane and copper cable electrical interfaces.

LPO: Linear pluggable optic is a technology used in optical transceivers that simplifies the design of pluggable optical modules by removing the traditional Digital Signal Processor (DSP) and Clock Data Recovery (CDR) chips. Instead, LPO utilizes a direct-drive linear approach where the signal path is considered linear, relying on the capabilities of the Application Specific Integrated Circuit (ASIC) in the host system (like a switch or Network Interface Card) to perform signal conditioning and equalization.

MCM: Multi chip module, a specialized electronic package where multiple integrated circuits (ICs), semiconductor dies or other discrete components are packaged onto a unifying substrate, facilitating their use as a single component (as though a larger IC).

Mid-board optics: an optical transceiver that is mounted on a PCBA away from the PCBA edge, close to a switch ASIC to reduce the amount of PCBA trace loss between an ASIC and the optical transceiver. This is in contrast to the common practice today of locating optical transceivers at the PCBA edge.

Midplane: Some backplanes are constructed with slots for connecting to devices on both sides and are referred to as midplanes.

MLSD: Maximum likelihood sequence detection is a mathematical algorithm to extract useful data out of a noisy data stream.

MR: Medium reach. CEI MR specifies chip-to-chip electrical interfaces.

NG: Next generation.

NRZ (PAM2): Non return to zero, a binary code in which 1s are represented by one significant condition (usually a positive voltage) and 0s are represented by some other significant condition (usually a negative voltage), with no other neutral or rest condition.

NPC: Near-package copper. NPC uses a copper cable to bring the front panel signal to a location close to the host silicon to minimize the host PCB losses. It reduces PCB losses by bringing the signals to a connector on the PCB close to the ASIC whereas CPC (Co-packaged copper) brings the signal to a connector on the ASIC package

NPO: Near-package optics. Similar to CPO (Co-packaged optics) and NPC (Near-package copper), NPO is an electrical to optical device intended to be mounted on the host PCB at a location adjacent to the host silicon to minimize host PCB traces to minimize electrical signaling requirements.

OE: Optical engine.

O-to-E and **E-to-O**: Optical to electrical interface and Electrical to optical interface, a component that converts an optical signal to an electrical signal or vice versa.

PAM: Pulse amplitude modulation, a form of signal modulation where the message information is encoded in the amplitude of a series of signal pulses. For optical links it refers to intensity modulation.

PAM4: Pulse amplitude modulation-4 is a two-bit modulation that takes two bits at a time and maps the signal amplitude to one of four possible levels.

PAM6: A digital signal modulation scheme that encodes information by varying the amplitude of a pulse across six distinct voltage levels. Each of these six levels can represent approximately 2.5 bits of data.

PAM8 A digital signal modulation scheme that encodes information by varying the amplitude of a pulse across eight distinct voltage levels. Each of these eight levels can represent 3 bits of data.

PCB/PCBA: Printed circuit board (PCB) assembly, an assembly of electrical components built on a rigid glass-reinforced epoxy-based board.

Repeater: A low-latency electronic device that receives a signal and retransmits it. Repeaters are used to extend transmissions so that the signal can cover a longer distance. Besides signal equalization, clock and data recovery (CDR) functions could be also added to remove jitter from received signals effectively.

RoCE: RDMA over Converged Ethernet (CE) is a network protocol which allows remote direct memory access (RDMA) over an Ethernet network.

RS: Reed Solomon FEC coding is a type of block code. Block codes work on fixed-size blocks (packets) of bits or symbols of predetermined size. It can detect and correct multiple random and burst errors.

RTLR: (Retimed Transmit, Linear Receive) also generically referred to as Linear receive optic (LRO), is a type of optical transceiver technology used primarily in high-speed data center and networking applications, especially within AI clusters. The RTLR naming convention is within OIF. LRO is characterized by the presence of a Digital Signal Processor (DSP) solely on the transmit path, while the receive path operates with a linear, non-retimed architecture. This differs from fully retimed optical modules that utilize DSPs on both transmit and receive paths, and also from Linear Pluggable Optics (LPO) that eliminate DSPs entirely.

Scale out: also known as horizontal scaling, refers to the process of increasing capacity and performance by adding more individual machines or nodes to a distributed system.

Scale up: also known as vertical scaling, refers to the process of increasing the capacity or performance of a single server or system within an AI data center by adding more resources.

SDO: standard development organizations.

SerDes: A Serializer/Deserializer is a pair of functional blocks commonly used in high-speed communications to transfer data over a relatively low number of lanes.

SFP: Small form-factor pluggable connector is a modular, hot-pluggable interface used in networking devices to connect to various types of fiber optic or copper cables. It uses a PCB card edge interface.

SI: Signal integrity is a set of measures of the quality of an electrical signal.

SNDR: Signal-to-noise-and-distortion ratio is a measurement of the purity of a signal.

SNR: Signal-to-noise ratio.

TBD: To be determined.

Tbps: Terabits per second. The throughput or data rate of a port or piece of equipment. Tbps is 1×10^{12} bits per second.

TCM: Trellis coded modulation is a technique that combines convolutional coding and modulation to improve data transmission efficiency over bandwidth-limited channels, like telephone lines. It achieves this by intelligently integrating the encoding and modulation processes, increasing the distance between signal points in the constellation to enhance error correction without expanding bandwidth.

TME: Test and measurement equipment.

Twinax copper cable: A type of copper cable similar to coaxial cable, but with two inner conductors instead of one.

UCIe: Universal chiplet interconnect express is an open industry standard that defines the interconnect between chiplets, or small component dies, within a single package.

VLC: Vertical line card. A new line card design in which vertical I/O connectors and ASIC are mounted side by side, reducing the signal trace distance.

VSR: Very short reach. CEI VSR specifies chip-to-module electrical interfaces.

XSR: Extra short reach. CEI XSR specifies die-to-optical engine (D2OE) and die-to-die (D2D) electrical interfaces.

448G: A generic name for an expected technology enabling data rates (including overhead) of approximately 448 Gbps per lane.

† Some definitions include content from www.wikipedia.com

1 Executive Summary

In the past 25+ years, the OIF has supported the communications industry by generating implementation agreements such as CEI-56G and CEI-112G, which have been openly shared with other industry standards bodies. Currently, the OIF is developing CEI-224G implementation agreements. These implementation agreements (IAs) define the parameters and required performance levels necessary to support the development of cost- and power-effective broad industry ecosystems.

As the OIF anticipates the next generation of higher electrical data rates at 448 Gbps per lane, it is becoming apparent that new technological solutions will be required at many levels within future communication and compute systems. The objective of this CEI-448G framework document is to identify and define the hardware application spaces that could benefit from future CEI-448G OIF Implementation Agreements across multiple levels of hardware. Identifying and defining these application spaces will allow the OIF and others in the industry to have a common language and understanding as decisions are made to initiate new development projects.

The technical challenges of next-generation data rate systems are discussed, as well as test and interoperability issues that will need to be addressed for the various interconnection applications. Although some technical options are mentioned, it is not the scope of this document to define specific technical solutions for these applications or the priority with which the application spaces should be addressed.

As in the past, it is critical that the industry maintains interoperable interfaces for application spaces to enable cost-effective components, subsystem, and system development and deployment. This will ensure interoperable fiber, connectors, electrical interfaces, etc. Identification of the critical application interconnections is the first step to meeting this requirement. The goal of this document is to build consensus across the industry on the application spaces and motivate the initiation of collaborative discussions that are required to generate a broadly agreed set of project developments and objectives.

2 Introduction

2.1 Purpose

The OIF Next Generation Electrical Interconnect Framework identifies application spaces for next-generation systems and areas for future work by the OIF and other standards bodies. The scope of this document explores next-generation electrical interconnection interfaces at 448 Gbps per lane.

Interconnection interfaces in a typical system are needed for die-to-die, die-to-OE (optical engine), chip-to-module, chip-to-chip within a PCBA (printed circuit board assembly), between two PCBAs over a backplane/midplane or a copper cable, or between two chassis or racks. These interfaces may be unidirectional or bidirectional, optical or electrical, and may support a range of data rates.

For each application space, the IAs that follow from this framework should identify requirements to support interoperability across the various application spaces for optical and electrical links. They may include, but are not limited to:

- Relative cost considerations
- Electrical link reach and loss budgets
- Power consumption (including pJ/bit power efficiency)
- Link Latency
- Management of SI parameter settings
- Channel requirements and general characteristics
- Signal levels, supplemental error correction coding, and target BER
- Test and measurement methodologies

The Framework Document may recommend a number of follow-on sub-projects to address interoperability for specific application spaces.

2.2 AI network and system architecture

Interconnects are an essential part of AI infrastructure. Today's frontier AI models already exceed the memory capacity offered on a single accelerator, which requires that the model be split up or "shared" across multiple accelerators connected with a high-bandwidth interconnect. To maintain a high performance-cost balanced system, accelerators are typically connected hierarchically with three networks as shown in **Figure 1**. The first AI network hierarchy, known as the "scale-up" network, is usually local to the rack or a small number of racks, and enables tightly coupled communication between the accelerators, often using proprietary protocols such as NVLINK, Infinity Fabric, or inter chip interconnect (ICI) as well as open protocols such as UALink or Ethernet. Due to the command / response nature of scale-up communication, low latency and lossless performance are important for efficiency reasons. Small physical distances are desired. Due to the memory-to-memory purpose of scale-up and the small physical distances involved, low latency and lossless performance are important for efficiency reasons. To interconnect the accelerators at the larger cluster scale, a next level of hierarchy, known as the AI "scale-out" network, is used, with an order of magnitude lower bandwidth than scale-up. Scale-out requires multiple tiers of network switching to span the cluster. Examples of scale-out network protocols include InfiniBand and UEC. An end-to-end transport protocol such as RoCE (if using an Ethernet-based approach), paired with an appropriate congestion control mechanism, is commonly used to improve the reliability under load of the scale-out network. Management, control, and training data are fed to the cluster by connection to the system CPUs over a third separate "front-end" network, which may be thought of as analogous to a typical web-scale or cloud datacenter network.

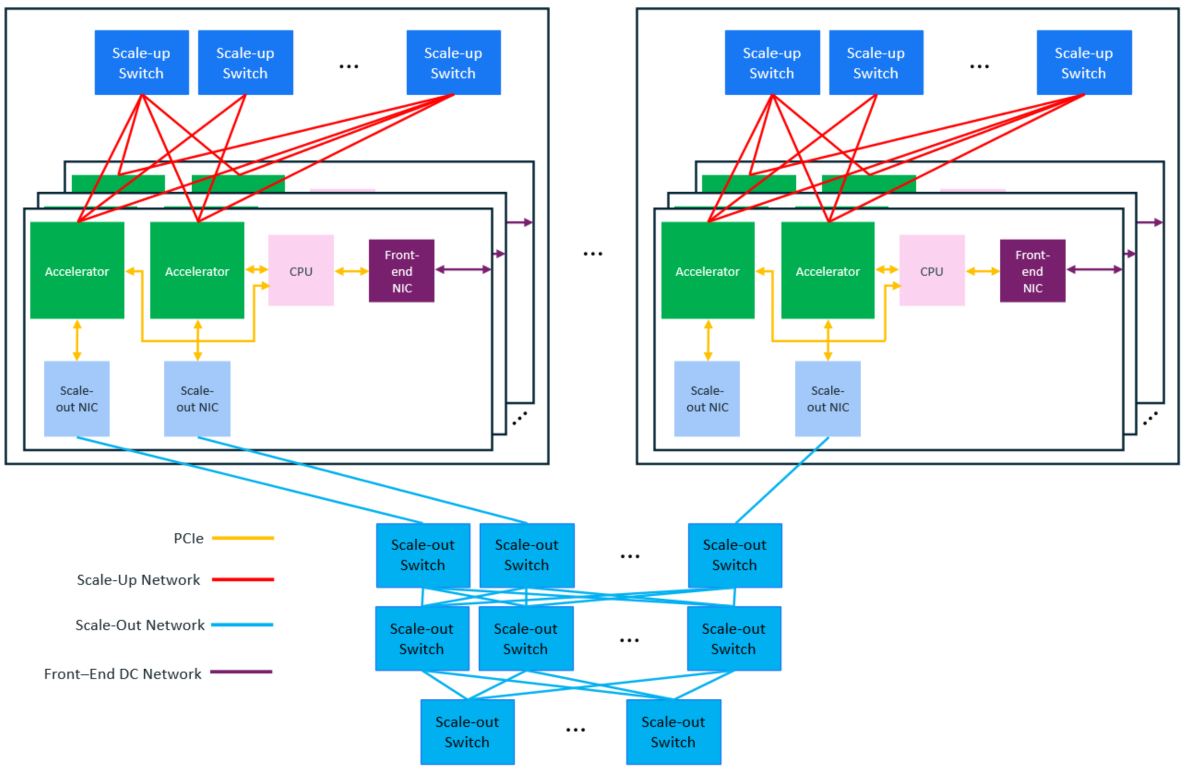

Figure 1 Block diagram of an AI network

The compute performance of a cluster is most heavily influenced by key metrics of the scale-up and scale-out networks, as outlined in **Table 1**. This is where the pressure for bandwidth is highest. As the scale-up network represents the majority of the interconnect bandwidth within the cluster and strongly influences performance, it is heavily optimized for cost, power, and reliability. The resulting common practice has been to use a copper interconnect where possible, often with a passive copper cabled backplane within the rack. Maximizing the number of accelerators to fit within a single scale-up domain in a power-efficient manner requires densification of the racks to reduce the physical channel distance and remain within the SerDes capability. Retimers, redrivers, or optics can also be used where necessary to extend reach, either to adjacent racks or within the rack for longer links.

Table 1 Key metrics of the scale-up and scale-out networks

	Scale-Up		Scale-Out	
	State of Art Today	Next Generation	State of Art Today	Next Generation
Number of Accelerator Nodes	~100	~1k	100k+	>>100k
Physical Size	Rack	Rack to Row	Datacenter	Datacenter
Network Properties	Lossless, Low Latency		Large scale	
Primary Interconnect Type	Electrical: Passive PCB / Twinax Backplane (within rack) Active Electrical Cables (AEC) (between adjacent racks)	Electrical: Twinax Backplane (within rack) Active Electrical Cables (AEC) (between adjacent racks) Optical (within row)	Optical	Optical

2.3 Comparative overview of SDOs for AI system interconnects

Many different standards development organizations (SDOs) are exploring the development of interconnect specifications to meet the rising needs of AI and HPC. Each SDO provides unique expertise and specific insights vital to establish a holistic view of the requirements, constraints, and tradeoffs to be considered. Extending OIF interfaces to 448 Gbps per lane signaling rates to meet the needs of AI and HPC systems will require tight collaboration amongst many SDOs including: UALink, UEC, SNIA, OCP, and IEEE 802.3.

2.3.1 OIF

The OIF Common Electrical I/O (CEI) IA (specification) provides a superset of foundational electrical specifications for high-speed, low latency communication between chips, modules and backplanes. The CEI specs continue to evolve technologically to provide the essential ingredients for modular, high-performance systems. CEI provides support for many different interconnects, ranging from MCM (die to die) to XSR/XSR+ (used with co-packaged optics and optical engines), VSR (used with pluggable front panel optics), MR (used for co-located chip-to-chip or chip-to-module connections), to LR (used for long distance chip-to-chip or chip-to-module connections). The OIF's ongoing 448G work is set to remain a driving force behind advancements across the industry ecosystem and influence next-generation interconnect standards. In addition, the OIF's EEI (Energy Efficient Interfaces) work is focused on addressing power optimized electrical and optical interfaces specifically for AI and HPC applications. Also related to support for AI and HPC applications is the OIF's Coherent Optical track addressing DCI (data center interconnect) links and Management track where CMIS (Common management information specification) interface extensions will be addressed for 448G separable electrical and optical link management requirements. Given the wide range of activities supporting next generation data rates, OIF will need to collaborate and liaison closely with the below SDOs.

2.3.2 Ultra Accelerator Link (UALink) Consortium

The UALink consortium develops open AI interconnect specifications optimized for direct load/store operations for low-latency accelerator-to-accelerator communication known as the "scale-up" network. The UALink 1.0 specification defines Transport Layer (TL), Data Link Layer (DL) and Physical Layer (PL) features optimized for the specialized, "scale up" network. The PL Technical Working Group (TWG) recently completed the 200 Gbps/lane signaling effort and is initiating the gathering of requirements and enhancements for the next signaling rate of ~400 Gbps/lane. Close collaboration with UALink will be required to capture their scale-up network requirements for OIF project starts.

2.3.3 Ultra Ethernet Consortium (UEC)

The Ultra Ethernet Consortium (UEC) is an open industry group focused on addressing the increased interest in high-performance compute (HPC) networking solutions based on Ethernet. The UEC specification is currently focused on the backend "scale-out" network (connections deployed across a "cluster") as well as considering opportunistic support for front end and scale-up networks. UEC enhances Ethernet's capabilities with purpose-built solutions for the unique demands of HPC and AI via the development of Transport Layer, Link Layer and Physical Layer, among others, extensions to Ethernet. The UEC Physical Layer group is actively surveying members to identify enhancements that will improve performance, reduce latency, and increase robustness at the 400 Gbps/lane signaling rate, as part of ongoing efforts to evolve the physical layer for future HPC and AI-optimized networks. Close collaboration with UEC will be required to capture their various interconnect requirements for OIF project starts.

2.3.4 SNIA

The Storage Networking Industry Association (SNIA) is another key standards development organization for interconnect and physical interfaces. In particular, the Small Form Factor (SFF) Technology Affiliate (TA) technical working group has a role defining cables, connectors, form factor sizes and housing dimensions, management interfaces, transceiver interfaces, electrical interfaces, and related technologies that are essential for next generation AI and HPC interconnects. These aspects will be critical elements in the increased signaling rate of front end, scale out and scale up networks at 400 Gbps/lane. Close collaboration with the SNIA TA will be required to define and ensure the cables, connectors, and other interface needs are captured for OIF project starts.

2.3.5 OCP

The Open Compute Project (OCP) is a collaborative community focused on redesigning hardware technology to efficiently support the growing demands of data center and cloud computing infrastructure through the creation of open standards. OCP has numerous projects related to AI and HPC use cases, including Open Systems for AI strategic initiative, having work areas in Open Cluster Designs for AI, and Short Reach Optical Interconnects, among others. Close collaboration with OCP will be required to ensure that the OIF defined interfaces support the high-speed, low latency data movements needed by their large-scale AI training and inference use cases.

2.3.6 IEEE 802.3

The IEEE 802.3 Working Group is responsible for developing and maintaining standards for Ethernet, the dominant technology used in wired local area networks (LANs). The Ethernet standard covers the media access control (MAC) and protocols for Ethernet operation, including the definition of interoperable physical layer specifications. The IEEE 802.3 New Ethernet Applications (NEA) ad hoc initiated an Ethernet for AI assessment to explore extensions to evolve Ethernet to support the needs of AI/ML and HPC networks. Close collaboration with the IEEE 802.3 will be required to ensure the harmonization of OIF project requirements with anticipated IEEE 802.3 400 Gbps/lane Ethernet physical layer developments.

2.4 Motivation of CEI-448G projects

With the growth of high-speed networks, next generation systems are being driven by the need to handle the increasing volume of data traffic. Recently Ethernet IEEE 802.3, SNIA SFF, UALink Consortium, Ultra Ethernet Consortium (UEC), Open Compute Project (OCP) and OIF CEI have kicked off their efforts targeting signaling beyond 224G per lane. So, what comes next beyond 224G for electrical interfaces over copper (Cu) channels? Will it be 448G?

To meet the next-generation system bandwidth requirement, Ethernet is aiming at aggregate interconnects and modules with 1.6TbE and up to 3.2TbE capacity and next generation switches targeting 200T capacity. Other industry organizations' roadmaps are also facing needs for these higher data rates as well. A temporary solution could be increasing the number of electrical lanes. However, this on its own will not be sufficient to meet bandwidth growth and will not provide a long-term solution for the best integration density and cost. The next generation systems are constrained by limits on power consumption, by limits on the size of a system, and by the need to provide a cost-effective solution. Higher signaling rate lanes would address these limitations. Therefore, 448G is critical to scale the device capacity and efficiency of networking and compute bandwidth.

3 Challenges and Potential Solution Spaces

3.1 AI scaling challenges and bottlenecks

Improvements in ASIC, memory, and packaging technologies have enabled generational increases in memory capacity per accelerator, which in turn drives a need to increase both scale-up and scale-out interconnect bandwidths. This ensures that accelerators spend the majority of their time performing computational work, rather than idling while waiting for communication.

To keep pace with the demands of frontier models, which continue to grow in model parameter size at approximately 4x per year, AI devices are being developed with increased memory, compute, and interconnect bandwidth. Today's leading devices already utilize the latest PHY technologies, based on CEI-224G standards for electrical interfaces, typically paired with low-loss media such as intra-system cabled host internal cables, cabled backplanes, or ultra-low-loss PCB materials.

Although interconnect bandwidth can be increased by adding more physical lanes, system-wide bottlenecks and congestion still prevent linear scaling. These include physical limitations in ASIC package size, PCB wiring, and electrical connector and cable size. Expanding the volume of these interconnects to accommodate more lanes is undesirable, as it would result in a lower-density system with fewer accelerators per scale-up domain or rack, working against the optimization of performance and cost. For this reason, it is essential for system designers to preserve the density of today's systems and interconnect solutions while scaling bandwidth through increased physical lane speed. Accommodations will need to be made at higher Nyquist frequencies, as crosstalk, reflections, inter-pair skew sensitivity, and material loss all increase. This likely requires the introduction of new approaches to minimize impairments, such as copper cabling terminated directly to the ASIC package. This work will require close collaboration within the industry to ensure the development of robust, application-suitable, and interoperable solutions. **Table 4** in Section 4.12 summarizes a list of interface types either currently in-use or in contention for future use in AI focused applications.

3.2 Challenges of cost, power and electrical link reach

3.2.1 Constraints and complexities

The proliferation of AI workloads and the exponential growth of web-scale data centers have positioned high-speed electrical interconnects at the core of modern infrastructure. These interconnects are now central to enabling the performance and scale required for emerging applications. However, the transition from 224G to 448G is more than an incremental step; it represents a paradigm shift in system design. This transition demands innovation across every facet of the technology stack, from materials and packaging to system architecture and thermal management.

Historically, incremental advances in bandwidth were achieved by optimizing individual domains, such as electrical and optical signaling, thermal, or mechanical design. However, at 448G, this approach is no longer sufficient. Achieving comparable reach at these higher speeds is exceptionally challenging due to several factors. Diminishing margins resulting from increased data rates drastically reduce signal integrity and power delivery headroom, making systems far more sensitive to design trade-offs. At the same time, the scale of AI and hyperscale deployments significantly raises the probability of link failure, necessitating more stringent reliability standards. Furthermore, shorter product cycles and rapid technological advancements now demand that new materials, modulation schemes, packaging, connectors and cooling solutions be developed and brought to market at an unprecedented pace.

Table 2 Historical evolution of the CEI-LR projects

OIF CEI projects	CEI-56G-LR	CEI-112G-LR	CEI-224G-LR	CEI-448G-LR
Timeline	2014-2017	2017-2021	2021-	2026-
Ethernet rate	50/100/200G	100/200/400G	200/400/800/1600G	400/800/1600G/3200G
Switch capacity	12.5T	25T/50T	50T/100T	100T/200T
Per-lane data rate	56 Gbps	112 Gbps	224 Gbps	448 Gbps
Modulation	PAM4	PAM4	PAM4	TBD
Insertion loss	30 dB at 14 GHz ball-ball	28 dB at 28 GHz ball-ball	40 dB at 56 GHz bump-bump	TBD
Reach objectives	3m copper cable	2m copper cable	1m copper cable	TBD
Pre-FEC BER target	1e-4	1e-4	1e-4	TBD
SerDes architecture	Analog/DSP	Analog/DSP	DSP	TBD

3.2.1.1 Electrical channel limitations

The increased Nyquist frequency for 448G poses significant challenges to bandwidth-limited copper interconnects. The industry has successfully developed three generations of high-speed interconnects-56G, 112G, and 224G—using PAM4 modulation as shown in **Table *2***. However, the feasibility of extending PAM4 to 448G is now under scrutiny. A primary challenge is the difficulty in eliminating resonance-induced notches in the insertion loss curve above 90 GHz, which can significantly degrade signal quality and limit the achievable link reach. Such resonance is primarily induced by the connectors in the link, particularly by the stub of the connector pin. The pace and timing of technological innovation in the connector industry will be crucial in determining the feasibility of PAM4 for 448G.

To address the escalating demands of higher data rates, higher-order modulation schemes such as PAM6 and PAM8 are being considered. These schemes offer the advantage of lowering the required Nyquist frequency, which can help alleviate some channel bandwidth constraints. Nevertheless, they introduce new challenges: a reduced signal-to-noise ratio and narrower receiver skew tolerance windows. This heightened sensitivity to channel impairments places great demands on the quality of cables and connectors, as well as on overall system engineering.

3.2.1.2 Power consumption: a critical constraint

The pursuit of longer reach in high-speed interconnects must be carefully balanced against its impact on interconnect power consumption. As highlighted in the OIF Next Generation CEI-224G Framework Document, bandwidth has increased 80-fold while interconnect power has risen 22-fold from 2010 to 2022, placing significant strain on the ability to deliver power efficiently to high-speed links. Effective power management at these data rates relies on minimizing SerDes power, optimizing channel performance, and reducing the total number of SerDes and repeaters in the signal path. For 448G, stringent signal integrity and loss requirements will likely necessitate DSP-based SerDes using advanced equalization techniques such as Maximum Likelihood Sequence Detection (MLSD), which can further increase SerDes power. Consequently, reducing channel loss—rather than simply adding repeaters—remains essential to maintain and further improve the downward trend in power per bit. On the other hand, with today's optical power dissipation, pJ/bit, simply replacing electrical links with optical links will expand the total power consumption significantly. Additionally, if optical and electrical interfaces have different modulation schemes the direct-drive approach is precluded, which results in increased power consumption.

3.2.1.3 Criticality of tail latency

For AI applications, especially those leveraging distributed training or inference across many nodes, latency—particularly tail latency, which refers to the delay experienced by the slowest packets or requests, typically measured at the 99th percentile or higher —can significantly hinder operations in several ways. During distributed training, all participating nodes must synchronize their results before moving to the next step, so if even one node experiences high latency in its data transfer, the entire group must wait, reducing overall efficiency and increasing training time. This can lead to resource underutilization, as some accelerators may sit idle waiting for synchronization, thereby diminishing hardware efficiency and increasing the cost per training run. This is important for AI links and especially important for scale-up links. As the size of the AI cluster grows, the likelihood of tail latency impacting performance increases, making it more difficult to scale up efficiently. Furthermore, for latency-sensitive inference tasks, high or unpredictable latency can delay responses, negatively affect user experience and reduce the reliability of real-time applications such as autonomous driving or live translation.

Controlling tail latency is therefore a critical design consideration in next-generation data center networks supporting AI and must be weighed alongside factors such as reach and power when evaluating trade-offs among different solutions, including modulation formats and FEC schemes.

3.2.2 Enabling important applications

Backplane and intra-rack connections continue to play a vital role, particularly in AI-driven scale-up networks where accelerators must be connected to switches within a rack. Electrical links throughout the rack are critical for their lower power, low latency and low cost attributes and therefore network operators rely heavily on maintaining these links at future rates. Co-packaged copper (CPC) cable solutions offer the potential for substantial loss improvements, but the practical reach achievable by SerDes—ideally without or with minimal repeaters—will determine whether copper interconnects remain viable for such applications, or if optical solutions become necessary even within the rack.

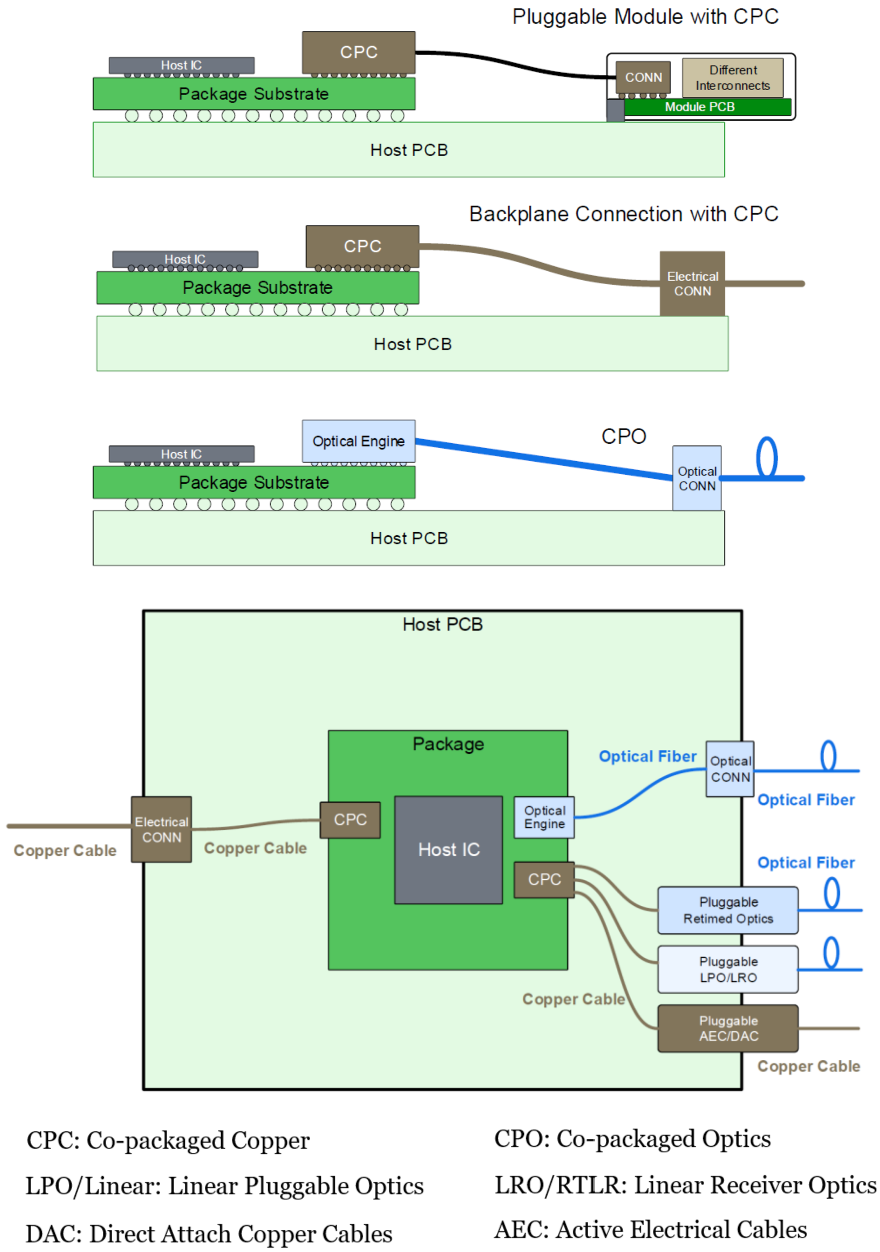

Figure 2 Leading applications for 448G

Enabling pluggable modules remains crucial for operational flexibility and ease of maintenance. System architectures, such as Vertical Line Card (VLC) and co-packaged copper cable, are under active evaluation to minimize insertion loss from the silicon to pluggable modules and to optimize the balance between performance, efficiency, and operational simplicity. The PCB edge connection, which has been in use since the introduction of SFP in 2001 and enabled many generations of pluggable copper cable and optics, including 1.6T modules with 8 lanes at 224G PAM4 signaling—is now being challenged to address signal integrity issues at 448G. Consequently, the industry is evaluating the need and requirements for high-density connectors through various initiatives, such as the OIF High-Density Connector Project, for 448G and future copper and optical interconnects.

3.2.2.1 The promise and challenges of co-packaged optics

Co-packaged optics (CPO) offers an effective way to minimize electrical channel loss, making it an attractive option for next-generation data center architectures. It presents an opportunity to continue utilizing techniques that reduce SerDes numbers, such as linear optics at 448G. However, its widespread deployment is currently constrained by factors such as ecosystem immaturity, integration complexity, concerns around reliability and serviceability, and higher initial costs. Additionally, existing investments in pluggable optics infrastructure and the need for proven long-term operational data contribute to a cautious approach among operators. As these challenges are addressed and the technology matures—particularly considering the challenges of overcoming electrical channel bandwidth limitations—co-packaged optics (and near-packaged optics) may see broader adoption in data center environments at the 448G generation.

3.2.2.2 Cooling

Another significant contributor to overall power consumption in high-speed interconnect systems is cooling. As power densities rise—especially with high-radix switch designs and 448G interconnects—traditional air cooling is reaching its practical limits. Robust thermal management is essential for reliable operation, consistent performance, and protecting components from overheating.

To address these challenges, liquid cooling—which offers much higher heat removal efficiency than air—may become the mainstream solution for 448G systems. In closed-loop setups, coolant circulates through cold plates or heat exchangers attached directly to heat-generating components such as ASICs or optical modules, transferring heat to a secondary medium like air or facility water before final dissipation. **Figure 3** illustrates two abstract examples of how liquid cooling could be utilized. Additionally, ongoing industry discussions and conceptual demonstrations are exploring in-module liquid cooling for future 3.2T and 6.4T coherent modules, which are expected to dissipate over 100W and may require these advanced cooling solutions.

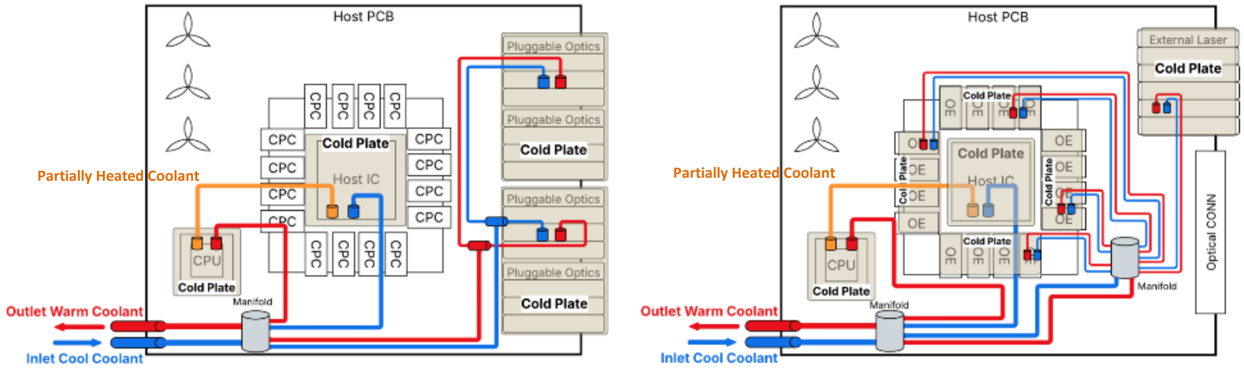

Co-packaged Copper Cable with Pluggable Optics
Liquid Cooled Host IC, CPU and Pluggable Optics

Co-packaged Optics with External Laser Modules
Liquid Cooled Host IC, CPU, Optical Engines, and External Laser Modules

Figure 3 Two conceptual approaches to liquid cooling

While immersion cooling, which submerges hardware in a thermally conductive dielectric liquid, is unlikely to be widely adopted at 448G, it offers even greater heat removal potential and may become relevant as power densities continue to climb in future systems.

Deploying advanced cooling solutions brings new system-level challenges. Effective air cooling depends on optimized airflow management and high-performance fans, but increased fan power also impacts total energy consumption. Liquid cooling systems require reliable fluid distribution, precise flow control, continuous monitoring, and robust leak detection to protect sensitive electronics and ensure uptime.

3.2.3 Summary: a paradigm shift

The transition to 448G electrical interconnects marks a fundamental shift in data center and AI system design, introducing unprecedented challenges across cost, power, reach, reliability, and operational complexity. Realizing the full potential of 448G will require a holistic, cross-disciplinary approach—addressing technical, operational, and economic constraints in parallel. Collaboration across the ecosystem, from component suppliers to system integrators and data center operators, will be essential to deliver scalable, efficient, and reliable high-speed interconnect solutions for the next generation of digital infrastructure.

3.3 Challenges of management of SI parameter settings

Ensuring interoperability across chip-to-chip (C2C/MR) and chip-to-module (C2M/VSR) electrical interfaces presents a significant challenge in standardizing the management of signal-integrity (SI) parameter settings. This is particularly crucial for enabling "plug and play" functionality for copper and optical modules (both front-panel pluggable optics and co-packaged optics) from various suppliers and with different reaches. Each successive higher-speed generation of C2C/MR and C2M/VSR has introduced new configuration methods, including programmed SI values, autonomous adaptive determination of settings, a combination of both, and, most recently, in-band link training for modifying electrical transmitter characteristics—a technique used by the industry for generations of direct attach copper (DAC) cables.

Host integration of a module becomes further challenging due to the necessity of a defined startup sequence, ensuring harmonized expectations regarding the time order of events between the host and the module. The industry addressed this by developing the Common Management Interface Specification (CMIS), which defines, among other things, a startup sequence for modules. The latest CMIS revision, 5.3, was published by OIF on September 4, 2024, and is available at https://www.oiforum.com/wp-content/uploads/OIF-CMIS-05.3.pdf.

Let's examine how these challenges have evolved for modules over the years:

For 10G-NRZ-lane C2M/VSR: The host programmed the module output equalization.

- This required the optical receiver to support electrical-output equalization, known as de-emphasis.
- Programming values ranged from 4.4 to 7 dB for Ethernet.

For 25G-NRZ-lane C2M/VSR: The host programmed the module input equalization.

- This required the optical transmitter to support electrical-input equalization, known as Continuous Time Linear Equalization (CTLE).
- Programming values ranged from 1 to 9 dB in 1 dB steps for Ethernet.

For 50G-PAM4-lane C2M/VSR: The module featured adaptive input equalization.

- This required the optical transmitter to support autonomous adaptive CTLE.
- A defined startup sequence was necessary to establish the expected time order of events between the host and the module.
- This sequence addressed when the module and host would receive valid electrical input signals for adaptation.

For 100G-PAM4-lane C2M/VSR: The host programmed module output characteristics and supported adaptive input equalization.

- This required the optical transmitter to support autonomous adaptive CTLE.
- The optical receiver needed to support two programmed output configurations, simply referred to as “short” and “long,” corresponding to whether the host has short (lower loss) or long (higher loss) channels/traces.
- IEEE Std 802.3ck™-2022 defined AUI-S C2M/VSR and AUI-L C2M/VSR for Ethernet.
- CMIS-VCS specification and CMIS-LT specification were introduced to provide greater ability for the host to control and train optical modules such as LPO, CPO/NPO, and RTLR-based modules.

For 200G-PAM4-lane C2M/VSR: The host-to-module electrical interface utilizes in-band link training, similar to that used by copper cable, with the addition of a ready-to-send signal to manage the overall link bring-up.

For the 400G-lane generation, the primary challenge will be meeting the required BER for the target application, or, stated differently, achieving the minimum mean time needed to retransmit a frame/packet for the target application.

3.4 Challenges of channel requirements and characteristics

The importance of the electrical channels in network and compute architectures, especially to support AI and high-performance applications, cannot be denied. So, it is critical to understand the characteristics of the interconnect and interfaces to plot possible solutions for future data rates.

As one considers the opportunity for doubling the electrical data rates, a number of practical considerations emerge. Within comparable passive interconnect technology, increasing the operating data rate leads to degradation in the Signal-to-Noise Ratio (SNR). The amplitude of the signal component will decrease due to the increased conductive and dielectric losses of the medium. Due to shorter unit intervals at a higher rate, combined with inevitable mechanical feature transitions in the channel, the noise component due to reflections and crosstalk will increase and be harder to mitigate. Moreover, the accompanying need for increased I/O density further exacerbates the concern for cross talk due to trace and port proximity. Finally, power consumption has become a critical parameter for network operators. Hence, channel definitions will need to carefully balance the usage of repeaters and high-performance equalization schemes vs. latency and power consumption. Bounds on the power envelope and active component maximum feature set will drive improved channel performance. Where will all of this come from?

Traditional Long Reach chip-to-chip implementations can connect devices from different vendors together across a variety of mediums with a range of different interconnects. As we embark on the next data rate, it's important to assess the capabilities and limitations of existing form factors, topology and manufacturing techniques. Finding the performance bottleneck in the channel early can help stimulate the conversation to enable new form factors and technology that can help focus on the solution space for the next generation of implementation agreements. The areas marked in BLUE denote key transition points between medium and vendors that are often the source of critical discontinuities and crosstalk that degrade channel performance. The areas marked in ORANGE below represent areas of industry consensus for channel loss budget allocations or optimized characteristic impedance. During the 224G IA development there was a strong pull towards 92.5-ohm interconnects and that is expected to continue for 448G efforts.

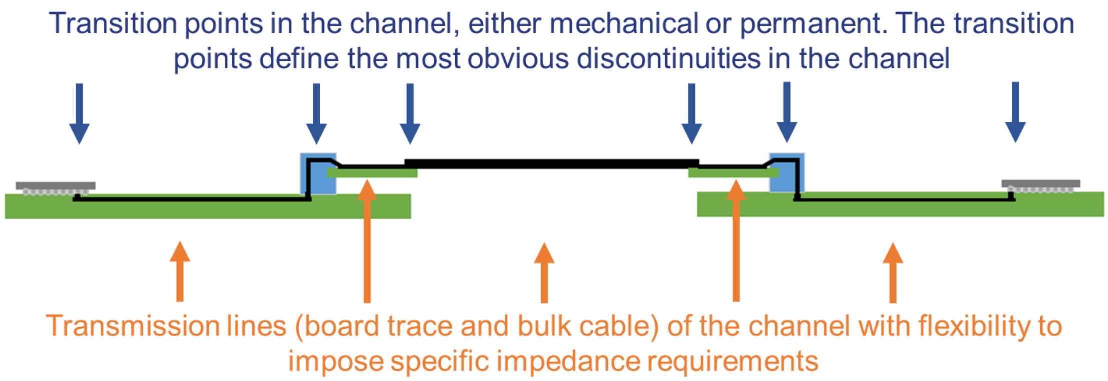

Figure 4 Illustration of a CEI-LR type implementation, highlighting critical channel aspects

When considering physical reach of an electrical channel, current materials that make up the medium significantly contribute to the total signal attenuation. This is agnostic of applications from short reaches such as die-to-die electrical links in a multi-chip module or an optical engine, to the other end of a spectrum when considering longer reaches such as backplane and copper cable applications. The chip package and printed circuit board (PCB) materials (dielectric and conductor) have an impact on the total loss. Advancements are being made in these materials to support higher data rates, but still other media will likely need to be considered such as twinax cable for internal applications such as “cabled-hosts” to preserve the channel reach. In addition, connectors will need to be improved for optimized high speed performance especially including bandwidth, insertion loss, return loss and cross-talk. As mentioned above, one of the effects that became clear during 112 and 224G IA developments was the impact that return loss could have on low loss channels such as cabled-hosts. Based on this learning, 448G applications (and channel metrics) will need to more heavily weight the impact of factors beyond insertion loss.

One of those channel metrics beyond insertion loss is intra-pair skew. Intra-pair skew can amplify jitter and inter-symbol interference ultimately degrading the eye diagram. In extreme cases, intra-pair skew can lead to noticeable reflections and resonances in the transmitted signal. The intra-pair skew is a component of mode conversion, converting differential signals into common-mode signals. It’s important to appreciate the difference in the resulting characteristics of intra-pair in PCB structures compared to twin-ax cable structures. It is shown in **Figure 5** the intra-pair skew over the frequency spectrum for different transmission media will exhibit different profiles. The end-to-end channel requirement could be a composite signal integrity metric (i.e. Channel Operating Margin, or COM), while additional interconnect metrics could be standardized to isolate specific portions of the channel to control the impact of intra-pair skew and other electrical and physical phenomena that may exist at higher frequencies. As many systems and platforms deploy more cabling to meet physical reach expectations, it’s important that the topic of skew is driven to consensus regarding a standard interpretation of intra-pair skew and related system and interconnect requirements. The data shown in **Figure 5** is representative of the transmission media used in 224 Gbps channels and is not indicative of channel limitations for 448 Gbps. It is critically important that test fixtures and measurement methodologies extend beyond

67 GHz (used in 224 Gbps characterization) to support the challenges that may arise for 448 Gbps implementations.

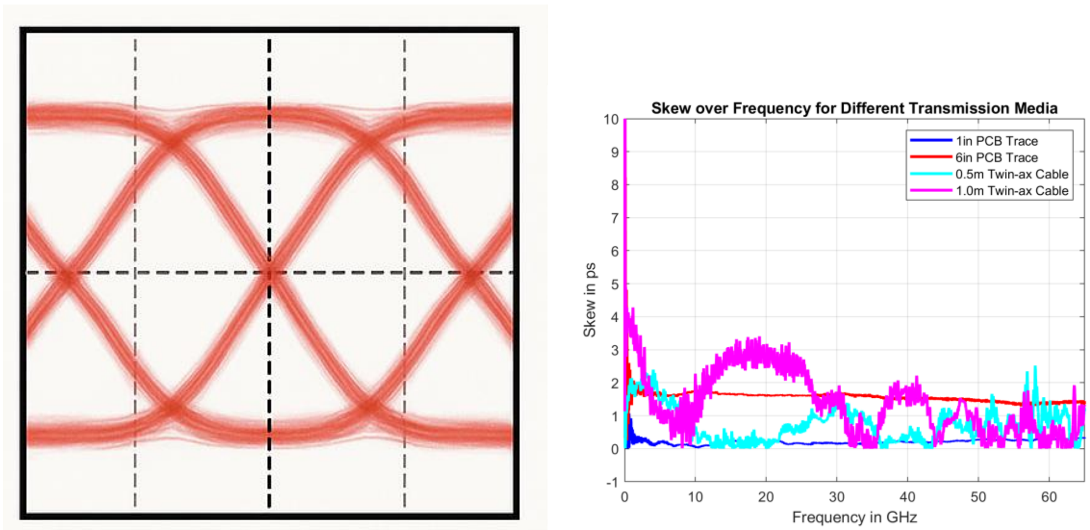

Figure 5 Intra-pair skew impacts an eye diagram and has characteristic attributes across the broadband frequency range

Another key factor to consider when discussing channel requirements is the channel bandwidth and potential modulation solution(s) that will be targeted for 448 Gbps. It is shown in **Figure 6** channel bandwidth tends to roll off around the 95 to 100 GHz range. This simulated data shows what may be achieved using a non-backwards compatible interconnect technology that is optimized for these next generation applications. Backwards compatibility is currently being discussed and may not be required since it brings a number of SI challenges and limited benefit.

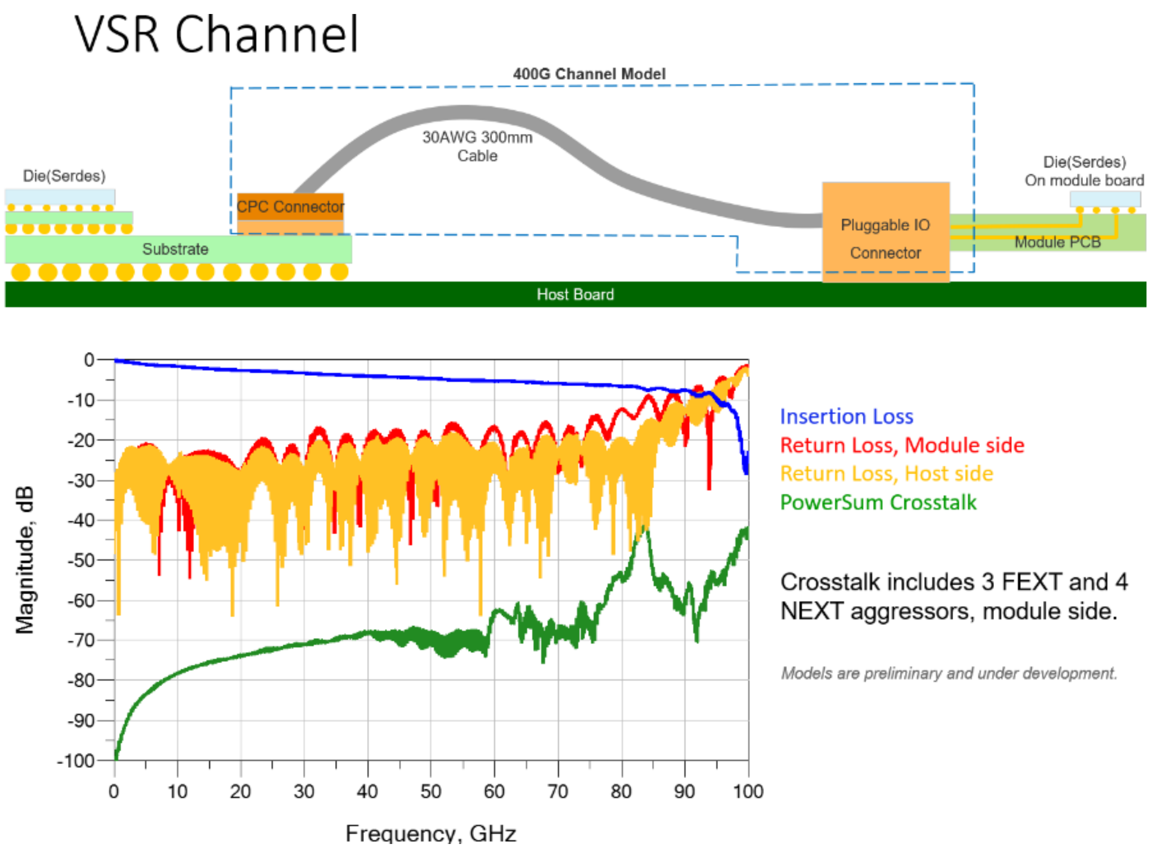

Figure 6 Illustration of a CEI-VSR type implementation, highlighting key interconnect performance characteristics

A modest increase in modulation complexity can reduce the signaling baud rate and significantly impact the quality and cost of the printed circuit board material that is required to support the channel reach or loss. In the case of 448 Gbps, PAM4 would result in a Nyquist frequency of 112 GHz, while PAM6, for example, would reduce Nyquist frequency down to 90 GHz, a significant improvement for loss performance (approximate values depending on overhead bandwidth required). On the other hand, the increased modulation level will also decrease the SNR for the end-to-end system. Moreover, this added modulation complexity may bring an engineering challenge to the chip design, but once done, may have a small incremental impact on the combined die functionality and cost. The topic of modulation choice recurs throughout the sections of this framework document as the industry considers the multiple tradeoffs for this critical topic. The impact of rack scale channels for advanced channel design/acquisition cost, and cost for advanced SerDes design/acquisition cost need to be both considered and balanced ("cost" in this usage includes acquisition cost, power, reliability, complexity, etc.). The same thought process must be considered before it is dismissed regarding modulation solutions for different reaches for example. While the benefit of having a common modulation solution across both optical and copper applications is obvious, again, the incremental cost of a 'dual mode' modulator where modulation 1 is optimized for lower loss chip to optical channels and modulation 2 is optimized for higher loss chip to chip copper channels, may be minor compared to the overall deployment costs when the volume of copper channels (exotic PCB material vs good PCB material) is considered. Similarly, a discussion may be required for the best ways to implement an optimized modulation conversion or 'gearbox' in applications if it is determined that electrical links and optical links cannot be optimized with the same modulation. Careful consideration of all channel implementations and modulation options is critical to ensure the 'rack scale' electrical channel architectures will be enabled including XSR (CPO), VSR (chip to module), MR (within a shelf) and LR (within a rack).

As industry considers IO port and circuit density increases driven by aggregate bandwidth roadmaps, great care will be required to ensure crosstalk is well managed in high density channel applications. Careful design optimization of connector and silicon package footprints as well as printed circuit board lithography processes are critical to mitigate the harmful efforts of these circuit transitions. In addition, increasing port densities can negatively impact channel reaches due to the need for noise isolation of traces. One key architecture to be evaluated for potential positive benefits is the termination of cabled host twinax cable solutions that directly terminate to chip packages rather than terminating to the host printed circuit boards near the chip package, Near Package Copper (NPC). This is referred to as "co-packaged copper" (CPC). This minimizes all the parasitic effects of the I/O connector to host PCB and host PCB to package transitions as well as reducing the package effects. A trade off that must be considered when evaluating increased port density and/or CPC is the escape routing challenges, and associated layer count increases due to the high channel densities that will be required. Additionally, to help enable interoperable channel performance, it may become necessary to further define the CPC mating interface dimensions and attachments for true multi-supplier interoperability, a technical and political hurdle. In **Figure 7** some of the traditional reference points (black arrows) are shown along with potential new points (red arrows) for consideration to support new electrical architectures that may require new demarcation points for interoperability and compliance.

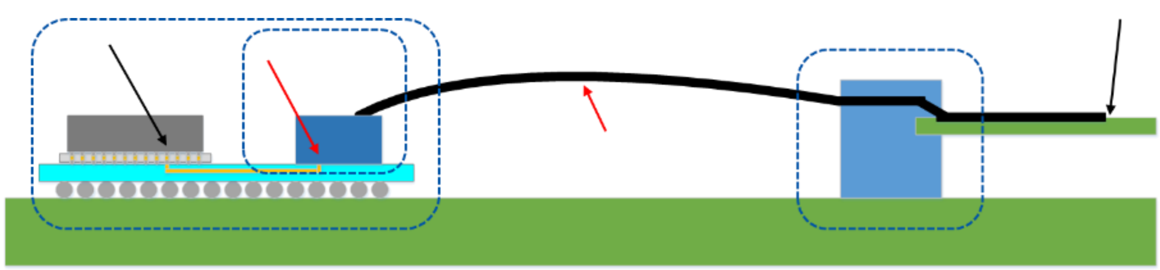

Figure 7 Illustration of possible reference points for 448G channel definition, highlighting practicality and interoperability

Due to the density and data rate requirements of 448G architectures and signal integrity, new interconnect form factors may be required and defined to enable interoperable channel specifications. As 448G requirements are being debated and analyzed, decisions may need to be made on new form factors that can support these requirements. As evidence of these discussions, consider OIF's 64 Channel HD Pluggable requirements project which is in the process of investigating and defining the electrical and mechanical performance requirements that would be required for next generation face plate pluggable electrical interfaces, while also considering requirements for co-packaged optics interface commonality as well. This discussion is also adding yet another dimension to the channel architecture requirements and that is thermal management. As data rate, channel density and reach all expand simultaneously, it is inevitable that we will also be considering methodologies and tradeoffs of various thermal management implementations that will conflict/tradeoff with electrical requirements. It is possible we may also consider discussions like this for other interfaces; CPC (mentioned above) and backplane as well.

Power consumption of a channel has become one of the highest priorities for network operators and is a critical consideration. Preservation of passive copper channels, where possible, enables the power reduction by not requiring short optical links, as long as the electrical transmitter and receiver can be realized in a power efficient manner. These silicon and channel tradeoffs must be carefully considered and debated to ensure the end user gets the best possible power efficiency. The industry has been very clear on the need to carefully design systems for power efficiency in light of the forecasted power envelope growth. As an example, whereas repeaters might have been sought to mitigate channel reach limitation, their use may not fit in the system power envelope. Therefore, the margin needs to be accommodated by a cleaner passive channel.

3.5 Challenges of modulation, equalization, target DER, and FEC/latency

CEI has adopted PAM4 as its modulation format since 56G, and PAM4 has served the CEI well for three generations, from 56G, to 112G, and now 224G for various reaches, from XSR, to XSR+, VSR, MR, and LR. Naturally, it is desirable if PAM4 can serve the CEI for its fourth generation considering modulation backward compatibility and optical compatibility for the overall system for power, latency, and cost optimization. On the other hand, whether PAM4 can support CEI-448G depends critically on the I/O link technology speed/bandwidth and performance advancements, including SerDes equalization, channel components including package, break-out, connector, PCB/cable, and FEC.

In order for SerDes to support 448G-PAM4, its analog front-end (AFE) bandwidth (BW), including TX AFE/PLL, RX AFE/CTLE and CDR bandwidth and speed both need to be doubled, while its jitter, noise impairments, needs to be reduced by 2x in terms of time, compared with 224G-PAM4. To achieve those advancements, two technologies are essential, one is the ADC/DAC (Digital to Analog Converter) DSP architecture and related integrated circuits, and another is the advanced process technologies (e.g., <= 2nm). While 448G-PAM4 SerDes test chips have not yet been broadly demonstrated, the recent demonstrations of 1.6T coherent optics where its electrical bandwidth has achieved 100 GHz, and with a signal rate of 200 GBd, is very encouraging data in supporting the 448G-PAM4 SerDes feasibility.

To increase the bandwidth of the package from the 224G-PAM4 Nyquist of 56 GHz, to 448G-PAM4 Nyquist of 112 GHz, package material, structure/stack-up (e.g., layers, core thickness), surface roughness, trace width/space, use of skip layer, and BGA pitch and capacitance, all need to be improved. A key parameter determining the package bandwidth is the BGA pitch. When a BGA pitch of 0.5 mm or less is used, along with other package characteristic parameters mentioned, a package bandwidth of 115 GHz can be achieved. 0.5 mm package pitch technology is available today. With the next gen advanced package technology, it is feasible to achieve the insertion loss at 448G-PAM4 Nyquist near that of the 448G-PAM4 (e.g., 6 dB for Class A, and 9 dB for Class B).

Increasing the channel bandwidth between chip packages BGA-BGA to 448G-PAM4 Nyquist of 112 GHz is needed to support 448G-PAM4. At present time, 448 Gbps channel bandwidth is limited to ~90 GHz as shown in **Figure 6**, largely due to the connector technology. Such limitation requires the investigation of other higher order PAMn modulations. The following **Table 3** summarizes the PAMn modulation order, related bit per symbol, symbol rate, unit interval (UI), Nyquist frequency, and SNR delta, for PAM4 to PAM16.

Table 3 PAMn modulation characteristics at 448 Gbps

Data Rate (Gbps)	448	448	448	448	448	448
PAMn levels	4	5	6	7	8	16
Bit per symbol[1]	2	2.32	2.58	2.81	3	4
Symbol rate (GS/s)[1]	224	192.94	173.31	159.58	149.33	112
UI (ps)	4.46	5.18	5.77	6.27	6.70	8.93
Nyquist Frequency (GHz)	112	95.47	86.66	79.79	74.67	56
SNR Delta[2]	0	-2.01	-3.62	-4.96	-6.11	-12.08

Note 1: Assume $\log_2()$ mapping of bits to PAM symbols
Note 2: For BER=1e-5

With the channel bandwidth limited to ~90 GHz, PAM4 will not work, and higher order PAMn, such as PAM6 or even PAM8 need to be considered. While PAM6 or PAM8 can provide solution in principle for the channels with bandwidth limited to ~ 90 GHz, it also creates undesired consequences, including:

- Mismatch with the IMDD optical PMD modulation which is currently heading for PAM4 modulation.
- 448G linear pluggable optics (LPO) will not be possible, in contrast to 112G-LPO and 224G-LPO due to the host and optics modulation mismatch.
- Backward compatibility with CEI-56G/112G/224G-PAM4.

If connector/channel bandwidth continues to be limited at ~90 GHz, the electrical I/Os at 448 Gbps, such as CEI-448G-VSR/MR/LR, PAM4 will not be feasible, and higher order PAMn such as PAM6, or PAM8 will need to consider to have a solution path or feasibility. If connector/channel bandwidth can be improved to ~112 GHz and beyond, electrical I/Os at 448 Gbps, PAM4 would be feasible for CEI-VSR/MR/LR. In either scenario, end-to-end channel loss and reach both need to be advanced to > 40 dB at Nyquist, and > 1 meter cable reach, in order to support meaningful use cases (e.g., scale-up cable reach within a rack).

To enable continued alignment of electrical and optical modulation schemes, another direction to investigate if PAM4 electrical channels are difficult to implement is the use of alternate modulations for optics such as PAM6 or PAM8. This would enable continued lower power architectures such as LPO and RTLR. Again, it is critical to balance all the tradeoffs for rack level power, density, and compute capacity.

Interestingly, for AI scale-up and scale-out, direct drive CPO could be an option to alleviate electrical connector/channel bandwidth and reach limitations, and have the advantage to scale to the next rate beyond 448G without practical reach limitation. Under this scenario, host 448G-PAM4 SerDes can directly drive and receive signals to and from 448G-PAM4 optical engine (OE) reside in the 448G CPO, providing power, shoreline density, and no electrical bandwidth/reach limitation, compared with pluggable optics. However, the challenge remains if the 448G CPO technology is ready to be widely deployed when the 448G market window opens, which is anticipated to be at 2027-2028.

Beyond baseband PAMn modulation investigation, alternative or more advanced signaling schemes have been proposed to send data faster at 448G such as:

- Wide-and-slow approaches: These involve using multiple parallel, lower-speed channels to achieve the overall high data rate. While simpler in terms of individual channel modulation, the increased number of components and lanes could become physically unwieldy and expensive for very high bandwidth AI systems.
- Bi-directional transmission: This involves transmitting data in both directions over the same wire, which can improve spectral efficiency. However, it introduces complexities in signal separation, echo cancellation, and crosstalk mitigation.
- Single-ended signaling: While straightforward, single-ended signaling can achieve 448G by transmitting signals in half rates through two wires. However, it is generally more susceptible to noise and interference compared to differential signaling, making it less ideal for high-speed, long-reach applications where signal integrity is crucial.
- Multi-toned approaches (like DMT - Discrete Multi-Tone): DMT uses multiple sub-carriers (tones) to transmit data, offering potential advantages in terms of spectral efficiency and robustness to channel impairments. However, DMT can be more complex to implement compared to PAM. Studies suggest DMT may have advantages over PAM for long-reach, 448G data communication.

Ultimately, the choice of modulation technique for 448G involves weighing the tradeoffs between data rate, spectral efficiency, power consumption, complexity, reach, cost and backward compatibility. While PAMn offers a viable solution for increasing data rates within a given bandwidth, the ongoing evolution of signaling technologies will continue to explore diverse modulation schemes and their optimal application in various scenarios.

If PAMn or even higher order modulation schemes are employed to reduce the signal bandwidth to match the harsh frequency response of the channel, advanced signal processing and advanced forward error correction (FEC) schemes would be considered for CEI-448G electrical links.

As the processing node improves, advanced digital signal processing (DSP) after an analog-to-digital converter (ADC) can provide cost-effective and strong equalization for high loss and noisy channels. As shown in **Figure 8** CEI-224G-LR reference receiver equalization block diagram, multi-tap feed forward equalization (FFE) and decision feedback equalization (DFE) can work together to cancel inter-symbol interference (ISI) caused by channel insertion loss and return loss. Floating taps have been considered as well to cancel reflections at far-end time intervals. Advanced detection schemes such as maximum likelihood sequence detection (MLSD) and its variants may be used to further improve the equalization capability of the receiver without significantly enhancing noise and crosstalk from aggressor channels.

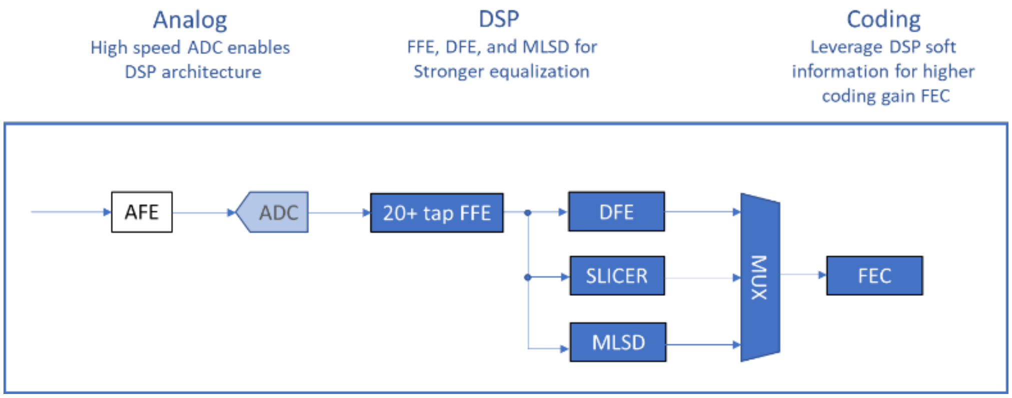

Figure 8 Block diagram of a CEI-224G-LR receiver

As data rates increase to 448 Gbps, these conventional techniques might not be sufficient to fight against channel loss and signal integrity challenges. More advanced equalization and signal processing methods, including noise and crosstalk cancellation, may be necessary.

It is anticipated that a pre-FEC Bit Error Ratio (BER) of 1e-4 and the burst error probability may not be achievable for a >= 40 dB channel at 448 Gbps. Consequently, a stronger FEC scheme than current KP4 FEC (i.e., RS (544, 514)) will likely be required for 448G electrical links.

Before looking into different FEC options, let's review the existing error correction architectures for 224G systems. Because of the SNR penalty of PAM4, FEC is assumed to be used in the system to achieve corrected CEI post-FEC BER of 1e-15 or better. There are three major types of FEC architectures in a multi-part link system as shown in **Figure 9**:

- End-to-End FEC: A single FEC encoder and decoder are shared across chip-to-chip, chip-to-module electrical links, and the optical link.
- Concatenated FEC: An inner code encoder and decoder are added at the module sides to provide additional protection for the optical link.
- Terminated FEC: Each electrical and optical link possesses its own dedicated encoder and decoder.

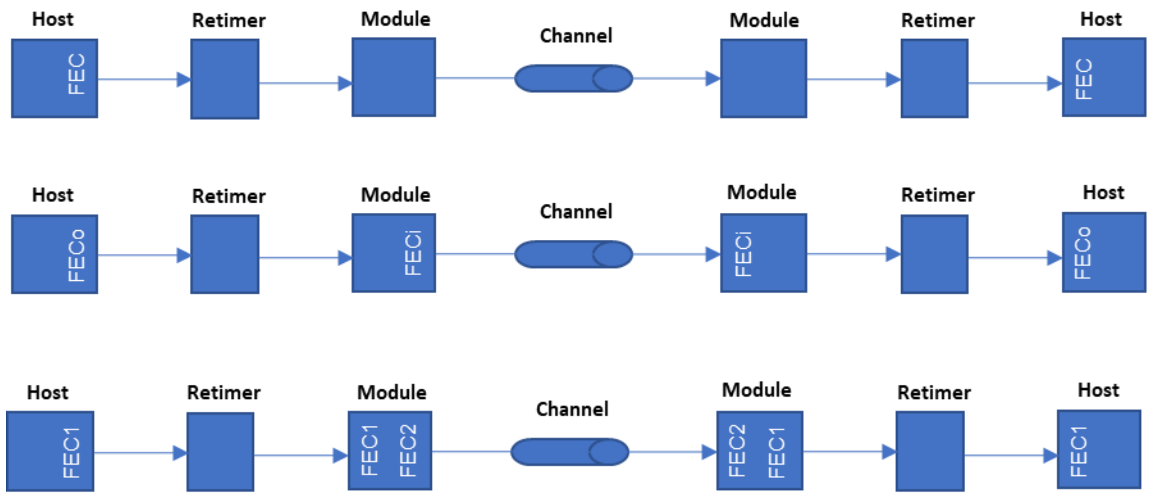

Figure 9 Existing 224G FEC schemes: end-end FEC (top), concatenated FEC (middle), terminated FEC (bottom)

These three FEC architectures are assumed to be viable for the 448G systems. Notably, there is no dedicated inner code encoder/decoder solely for 224G electrical links, but one may be needed for 448G electrical links.

A stronger FEC could be necessary for 448G electrical link to offset SNR penalty and provide performance margin. FEC can be enhanced by

- Increasing Redundancy: Introducing redundant bits allows the FEC decoder to correct more errors, improving signal reliability in noisy channels. However, this increases the overall transmission rate and consumes more bandwidth and increases the IL and reduces the SNR further, which is detrimental for bandwidth-limited electrical links.
- Expanding Modulation Constellation Space: Utilizing higher-order PAMn schemes, such as trellis coded modulation (TCM), can increase the number of bits transmitted per symbol without necessarily increasing the symbol rate, leading to higher spectral efficiency. Even though this makes the signal more susceptible to noise due to closer constellation points, with leveraging the redundant bit and reasonably selecting the mapping, the method could provide better SNR or lower latency.

However, stronger FEC for the electrical link to achieve better coding gains normally comes at the cost of coding overhead, coding latency and coding complexity. The trade-off of performance and cost for different applications needs to be carefully considered, especially for CEI-448G projects with low latency or/and low power application requirements.

3.6 Challenges of test and measurement

Testing and measurement methodologies for 448G electrical interfaces will face unprecedented challenges due to the extremely high data rates, complex modulation schemes, and stringent signal integrity requirements. Traditional test equipment and methodologies may no longer be sufficient. New approaches will be needed for:

- High-Bandwidth Probing and Fixturing: Accurate measurements at 448G require probes and fixtures with significantly higher bandwidth, lower insertion loss, and better return loss characteristics to minimize their impact on the signal.
- Module and host compliance fixtures will need to be developed for the corresponding form factors.
- Advanced Calibration Techniques: Calibration procedures will become more intricate to de-embed the effects of cables, connectors, and fixtures from the actual device under test (DUT) measurements.
- Complex Jitter and Noise Analysis: Characterizing jitter components (e.g., random, deterministic, periodic) and various noise sources (e.g., crosstalk, power supply noise) will require advanced analytical tools and more sophisticated measurement setups.
- PAMn Signal Generation and Analysis: Generating and analyzing PAMn signals at 448G requires highly linear and stable arbitrary waveform generators (AWGs) and oscilloscopes with sufficient bandwidth and vertical resolution. The increased number of levels in PAM6/PAM8 further complicates accurate signal analysis.
- BIT error rate test (BERT) will need to be developed to generate and analyze the appropriate PAMn patterns. Detailed analysis of the error types including burst length, distribution across the symbols needs to be reported. The BERT test sets need to support the appropriate baud rates used for 448G signaling. Error injection across PAMn symbols need to be supported.
- New PRBS patterns may need to be developed to adequately stress the channels and components required for 448G.
- FEC Performance Validation: Validating the performance of FEC schemes will involve generating and analyzing signals with controlled error rates and then verifying the correction capability of the FEC decoder. This requires specialized error analysis tools.
- Automated Compliance Testing: As complexity increases, automating compliance testing becomes essential to ensure efficient and repeatable validation of 448G components and systems against the defined IAs. This involves developing sophisticated test scripts and integrated testbeds.
- Thermal and Power Characterization: Measuring and analyzing the thermal performance and power consumption of 448G components under various operating conditions will be crucial, given the emphasis on energy efficiency and thermal management.
- Availability of components such as FPGAs, gearboxes & retimers –key elements of T&M – able to support the requirements of 448G multi-channel signal analysis and data generation.

Collaboration between the OIF, test and measurement equipment (TME) vendors, and component manufacturers will be vital to develop and standardize these new test methodologies and tools, ensuring consistent and reliable validation of 448G compliant products.

3.7 Summary

The transition to 448G electrical interconnects presents a multitude of significant technical challenges across various domains, including signal integrity, power efficiency, thermal management, and testability. The exponential growth of AI and HPC applications is driving an urgent need for higher bandwidth and lower latency, pushing the limits of current technologies. Addressing these challenges will require a holistic and collaborative approach, involving innovations in materials, packaging, modulation schemes, equalization techniques, and error correction. The OIF's framework document serves as a critical first step in defining the application spaces and initiating the necessary industry-wide discussions and development projects to realize interoperable, cost-effective, and robust 448G solutions for the next generation of digital infrastructure. The decisions made regarding modulation, FEC, power optimization, and test methodologies will collectively shape the future of high-speed electrical interconnects.

4 Interconnect Applications

Figure 10 illustrates the various interconnect application spaces that will be considered for 448G electrical interfaces, ranging from very short distances within a package to longer reaches across racks. Each of these application spaces presents unique technical requirements and constraints that must be addressed by future OIF Implementation Agreements.

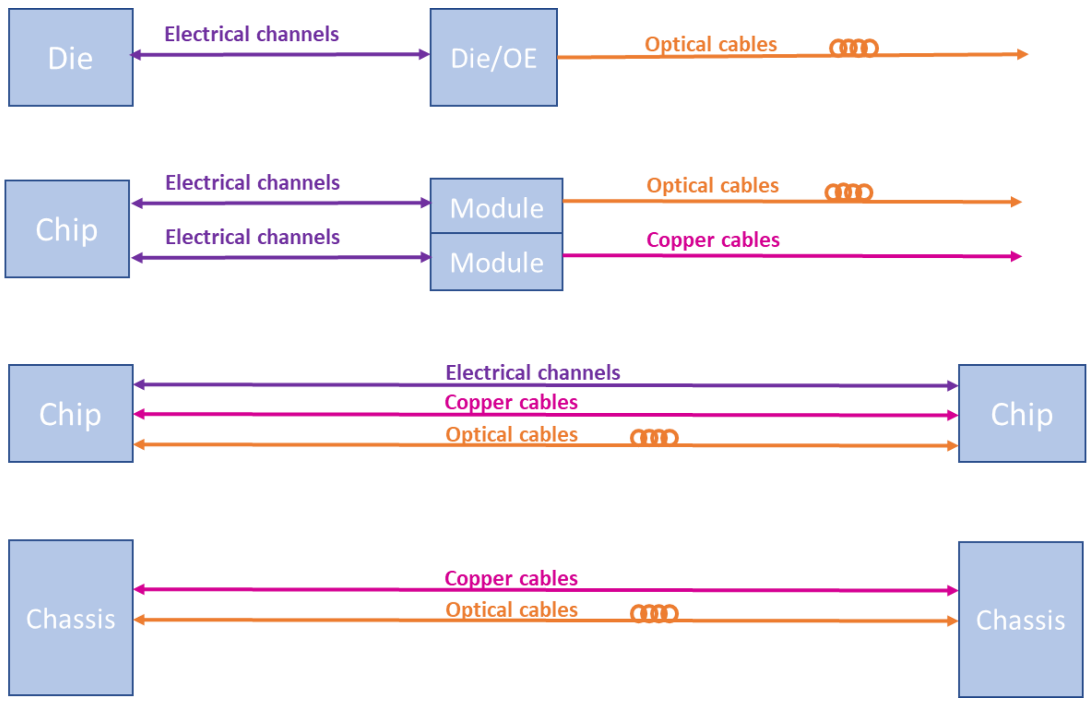

Figure 10 Interconnect application spaces

4.1 Die to die interconnect within a package

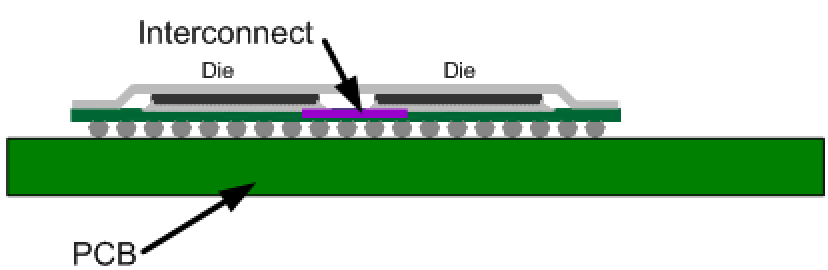

Figure 11 Die to die within an MCM interconnect application space

It may be necessary to use multiple dies within a multi-chip module (MCM) to achieve the industry's objectives. These co-packaged solutions can communicate with low power since the substrate provides a high-quality communication channel.

This short electrical link (typically sub-millimeter to a few millimeters) may allow for a much simpler interface and require less power than other longer standard electrical interfaces. For example, equalization is unlikely to be needed, and it may be possible to assume such short links are synchronous (single reference clock going to all chips), removing the need for a frequency tracking CDR. This is the XSR or XSR+ application space.

4.2 Die to optical engine within a package

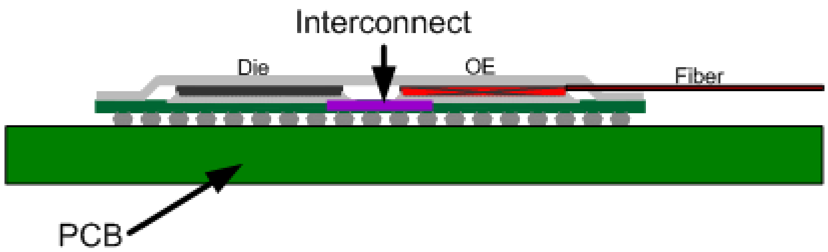

Figure 12 Die to optical engine MCM interconnect application space

This application space involves the electrical interconnection between a silicon die and an optical engine (OE) or co-packaged optic (CPO) within the same multi-chip module (MCM) or advanced package. The aim is to bring the optical conversion as close as possible to the high-speed silicon, minimizing electrical trace length and associated losses. These links are typically very short (on the order of a few millimeters to a few centimeters) and operate at very high densities. Similar to die-to-die interconnects, these interfaces benefit from the high-quality communication channel provided by the package substrate and the short reach, enabling low power consumption. This is a critical application space for the CEI-XSR or CEI-XSR+ specifications.

A key consideration for die-to-optical engine links is the potential for linear operation. If the optical link itself and the electrical link can align to the same modulation scheme, this could allow the optical engine to perform a linear conversion between the electrical and optical domains (O-to-E & E-to-O interfaces) without requiring complex Digital Signal Processing (DSP) for re-timing or re-shaping within the optical engine itself. This "linear" approach can lead to significant power and latency savings, which are highly desirable for AI and HPC applications. Evaluating a common modulation for both electrical and optical will be an important decision (as opposed to presuming the optical modulation).

4.3 SerDes to optical engine within an OE utilizing a slow-wide D2D interface from the ASIC

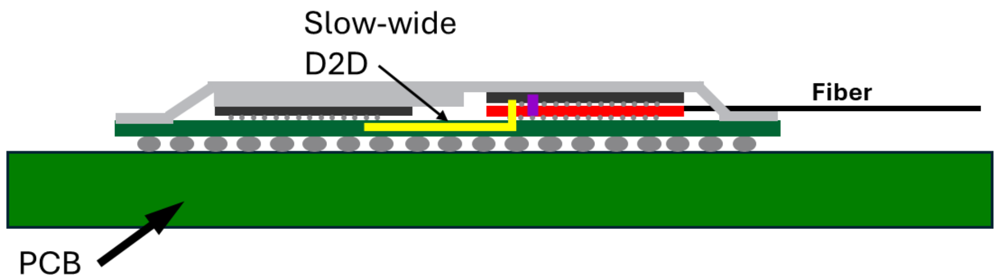

Figure 13 Optical engine die to PIC interconnect application space

This application space employs a slow-wide" D2D electrical interconnection between a silicon die and an optical engine (OE) on a shared packaging substrate. The high-speed SerDes is relocated from the die to the OE and the ASIC and OE are interconnected with a D2D interface. In many cases, the high-speed electrical channel inside the OE is extremely short vertical interconnection between the OE's electrical die and the OE's photonic die and ideal for "linear" implementations which are highly desirable for AI and HPC applications. Co-locating the SerDes with the OE can avoid any gearbox that would be needed between the electrical and optical modulation formats.

There are many D2D interfaces available which are energy efficient, and these interfaces support limited reaches from a few millimeters for some variants to a few centimeters for others. These D2D interfaces can be quite dense and can provide ASIC bandwidth escape edge densities well beyond 4 T/mm. UCIe and BoW are examples of these slow-wide D2D interfaces.

4.4 Chip to nearby optical engine

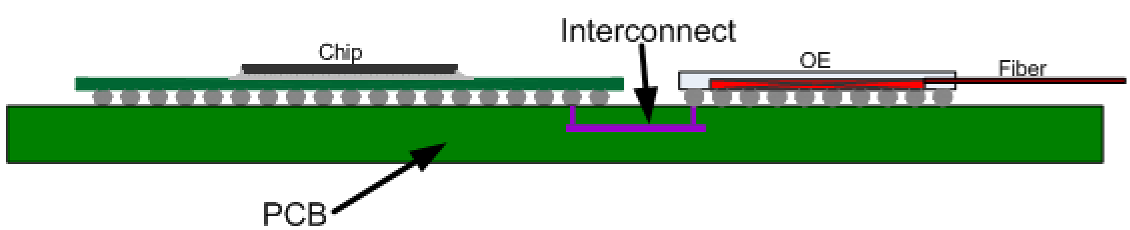

Figure 14 Chip to nearby OE interconnect application space

It may be useful to place an optical interface very close to the host chip (rather than placing the optical device within a host MCM due to heat restrictions of the optical components). In this case, a short electrical link through PCB trace or CPC cable is anticipated. Although this type of link will require more power than a link within a multi-chip module, the short reach of this channel would still imply that power could be saved. This electrical link for Near-Package Optics (NPO) or Near-Package Copper (NPC) interconnect could be CEI-XSR, CEI-XSR+ or linear interfaces.

The same modulation format may be appropriate for such links for backward compatibility and if the optical side uses an aligned format. However, advanced modulation formats beyond PAM4 or/and advanced FEC and equalization features are all possible solutions for the NPO and NPC interconnect.

4.5 Chip to module

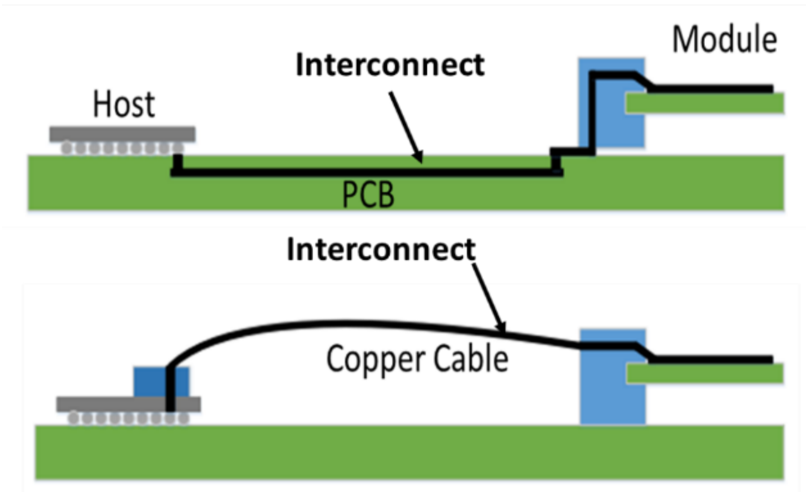

Figure 15 Chip to module interconnect application space

It is desirable to support pluggable modules at the front faceplate of the equipment. The traditional electrical link used to connect these pluggable modules was through a host PCB trace, a connector and a module PCB trace. Besides traditional PCB-based hosts (top), new cabled-host (bottom) implementations could be required to provide more port flexibility and link budget as shown in **Figure 15**. This is the VSR (Very Short Reach) application space.

At higher data rates placing retiming devices inside the pluggable module provides support for longer host traces but the inclusion of complex equalization features can overburden the limited power budgets of the pluggable module. Alternatively, a linear chip to optical engine (OE) interface could be considered to enable low power, low-cost, high-density form factor 448G serial optical modules in CPO, NPO and VSR applications.

Again, the same modulation format may be appropriate for such links for backward compatibility and if the optical side uses an aligned format. However, advanced modulation formats beyond PAM4 or/and advanced FEC and equalization features could be necessary for the chip to module interconnect to mitigate channel bandwidth limitation.

4.6 Chip to chip within PCBA

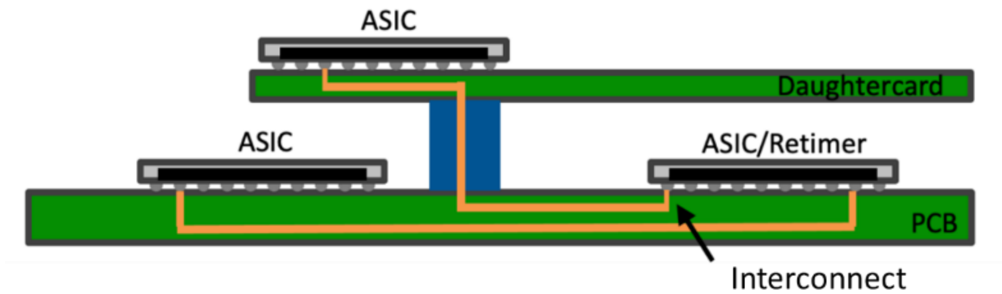

Figure 16 Chip to chip within PCBA interconnect application space

An interconnection interface may be needed between two chips on the same PCBA or on a daughter card or shorter mid-plane. By definition, this interface is through PCB trace or over the board cable and up to one connector. This is the MR (Medium Reach) application space.

Most chip to chip environments can save power if one can assume that both chips use the same power sources and the same reference clock, so that the signal noise sources are reduced in comparison to systems where the devices at each end of a channel are fully independent.

This interface would conventionally be electrical. It would, however, also be possible to use a combination of electrical and optical interfaces or even optical waveguides within the PCBA.

Advanced modulation formats beyond PAM4, FEC and equalization features are all possible solutions for the chip to chip interconnect. FEC may be a requirement to meet the BER – however the choice of the FEC must be considered carefully to address coding overhead, coding latency and power consumption concerns.

4.7 PCBA to PCBA across a backplane/midplane or a copper cable

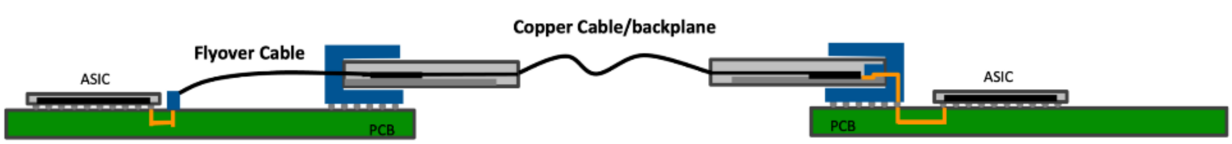

Figure 17 PCBA to PCBA across a backplane or a copper cable interconnect application space

An interconnection interface may be needed between two PCBAs that are located in the same chassis. Conventionally, this interface uses a backplane or midplane to connect the two PCBAs. However, as data rates increase, this can be achieved with copper cables that extend between the cards rather than through a backplane. These may be passive copper cables or active electrical cables. This is the LR (Long Reach) application space.

In addition, it may be appropriate to use advanced modulation formats beyond PAM4 in the link allowing for increased throughput density at a lower baud rate than PAM4. Stronger FEC may be a requirement to meet the BER. However, the choice of the FEC must be considered carefully to address coding overhead, coding latency and power consumption concerns.

4.8 Chassis to chassis within a rack

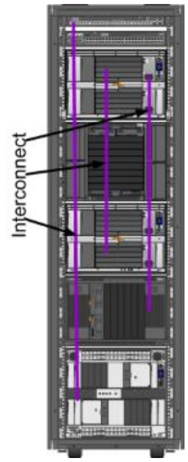

Figure 18 Chassis to chassis within the same rack interconnect application space

An interconnection interface may be needed between two chassis that are located in the same rack. This can be achieved with copper cables or optical cables. These may be passive copper cables, active electrical cables, or active optical cables. This is the LR (Long Reach) application space.

It may be appropriate to use advanced modulation formats beyond PAM4 in the link allowing for increased throughput density at a lower baud rate than PAM4. Stronger FEC may be a requirement to meet the BER. However, the choice of the FEC must be considered carefully to address coding overhead, coding latency and power consumption concerns.

4.9 Rack to rack side-by-side

Figure 19 Rack to rack side-by-side interconnect application space

An interconnection interface may be needed between two racks that are located side-by-side. This can be achieved with copper cables or optical cables. These may be passive copper cables, active electrical cables, or active optical cables. This is the LR (Long Reach) application space or beyond.

It may be appropriate to use advanced modulation formats beyond PAM4 in the link allowing for increased throughput density at a lower baud rate than PAM4. Stronger FEC may be a requirement to meet the BER. However, the choice of the FEC must be considered carefully to address coding overhead, coding latency and power consumption concerns.

4.10 Longer links

There is an increasing demand for longer reach links as AI clusters grow in size. For these links, optical interfaces are required. The OIF has been defining optical interfaces such as 400ZR and 800ZR that support these longer reach links. The OIF will continue to address the requirements for these longer reach optical links, and they may not require a specific 448G electrical interface for the host to module.

4.11 Retimed and linear interfaces

Linear Pluggable Optics (LPO) and Linear Receive Optics aka Retimed Transmit, Linear Receive (RTLR) are emerging solutions in high-speed networks, particularly driven by the escalating demands of data centers and artificial intelligence (AI) workloads. LPO, as the name suggests, utilizes a pluggable module that relies on the host system to handle signal conditioning and retiming, effectively removing the need for a Digital Signal Processor (DSP) chip within the transceiver module itself. This approach offers the advantages of lower power consumption, reduced cost, and potentially decreased latency, making it attractive for high-density, energy-efficient deployments like those found in AI clusters and cloud computing. The OIF CEI has developed CEI-112G-Linear IA and is currently working on the CEI-224G-Linear project.

However, LPO may face challenges with interoperability and link robustness due to its reliance on the host system. RTLR, also known as Half-Retimed Optics, presents a compromise by integrating a retimer in the transmit path while the receive signals are handled by the host system. This offers a balance between the benefits of LPO and the robustness of fully retimed optics, providing a cost-effective solution for scenarios where moderate signal integrity improvement is sufficient. The OIF EEI has developed EEI-112G-RTLR IA and continues working on the EEI-224G-RTLR project.

The industry has seen both retimed and linear electrical interfaces for high-speed signals. Retimed interfaces regenerate the signal at the receiver, which cleans up jitter and provides a robust link. Linear interfaces, on the other hand, simply amplify and equalize the signal without full regeneration. LPO and RTLR are typically lower power and lower latency, but are more sensitive to channel impairments. The choice between retimed, LPO and RTLR depends on the specific application, desired reach, power budget, and latency requirements. For 448G, both approaches may be considered for various interconnect applications, with a strong emphasis on power efficiency and minimizing latency for AI/HPC workloads.

4.12 Interconnect application summary

The above application spaces present a range of challenges and opportunities for 448G electrical interfaces. Each application has unique requirements in terms of reach, power, latency, and cost.

The OIF's framework will guide the development of specific Implementation Agreements to address these diverse needs, ensuring interoperability and facilitating the widespread adoption of next-generation high-speed interconnects.

Table 4 Interconnect applications

Interconnect Application	Distance Up To	Types of interfaces	Potential CEI projects
Die to Die in a Package	~25 mm	Electrical	XSR
Die to Optical Engine in a Package	~50 mm	Electrical	XSR/XSR+
Chip to nearby optical Engine	~150 mm	Electrical	XSR/XSR+
Chip to pluggable module	~250 mm	Electrical	VSR
Chip to chip within PCBA	~50 cm	Electrical or Optical	MR
PCBA to PCBA across a backplane/midplane/cable	~1 m	Electrical or Optical	LR
Chassis to Chassis within a rack	~1 m	Electrical or Optical	LR

5 Points of Interoperability

The Optical Internetworking Forum (OIF) promotes the development and deployment of interoperable networking solutions and services through the creation of Implementation Agreements (IAs) for optical networking products. It is therefore important for any next generation interconnects to consider the interoperability points to be defined in the agreement. The IA must also develop realistic measurement techniques for the defined interoperability test points.

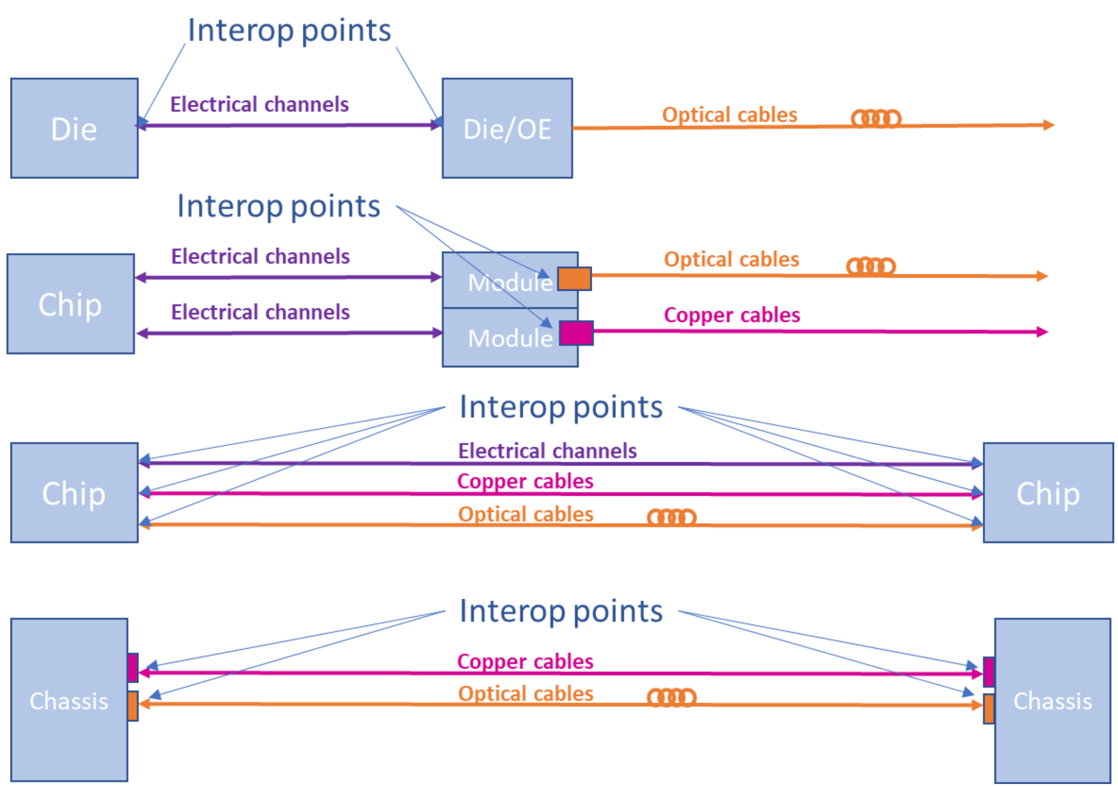

Figure 20 Interconnect application space showing points of interoperability

A next generation interconnect may be either electrical or optical. **Figure 20** illustrates the various points of interoperability that need to be considered when defining and developing solutions for high-speed interconnects, particularly at 448G. These points represent critical interfaces where different components and systems must seamlessly interact to ensure overall system functionality and performance.

The following points are crucial for next-generation systems:

- Chip-to-Package: This interface is critical for high-speed signaling between the silicon die and the package substrate. As data rates increase, managing signal integrity, power delivery, and thermal dissipation at this interface becomes increasingly challenging.

- Package-to-PCBA: This refers to the connection between the device package and the printed circuit board assembly. It includes the Ball Grid Array (BGA) or Land Grid Array (LGA) connections and the initial traces on the PCBA.
- PCBA-to-Connector: This interface involves the traces on the PCBA leading to a connector, which can be for a backplane, midplane, or direct cable attachment.
- Connector-to-Connector: This represents the connection between two connectors, such as those found on a backplane or when mating two separate cables.
- Cable/Backplane/Midplane: This encompasses the entire transmission medium, including copper traces on PCBs, twinaxial cables, or other structured cabling within the system or rack. The characteristics of this channel (insertion loss, return loss, crosstalk, skew) are paramount to signal integrity.
- Module-to-Host: This is the interface between a pluggable optical module (e.g., QSFP-DD, OSFP) and the host PCBA. It includes the module connector and the electrical traces within the module.
- Die-to-on-package: This is specific to co-packaged optics (CPO) and co-packaged copper (CPC), where the electrical interface between the host ASIC die and the optical engine or connector is very short and contained within the package.
- Die-to-near-package: This is specific to near-package optics (NPO) and near-package copper (NPC), where the electrical interface between the host ASIC package and the optical engine or connector is very short.
- Optical Interface: This refers to the optical connection between modules, including the optical fiber, connectors, and transceivers.

Ensuring robust interoperability at each of these points requires a comprehensive approach, including:

- Standardized Electrical Specifications: Defining clear electrical parameters (e.g., impedance, return loss, insertion loss, crosstalk limits, and jitter budget) for each interface.
- Common Mechanical Specifications: Ensuring physical compatibility of connectors, packages, and modules.
- Defined Management Interfaces: Utilizing specifications like CMIS for consistent configuration, control, and monitoring of optical modules and electrical interfaces.
- Interoperability Testing: Conducting rigorous testing with components from multiple vendors to validate compliance and ensure seamless operation.
- Reference Designs and Compliance Tools: Providing guidelines and tools to aid in the design and verification of compliant systems.

Collaboration among SDOs, industry forums, and individual companies is essential to achieve broad interoperability across these diverse points, especially as the complexity of high-speed interconnects continues to escalate at 448G and beyond.

6 448G Measurement Science

6.1 Intersection of Shannon-Hartley (theory) and physical systems

Decisions around modulation order (symbol rate) must leverage feedback driven by interconnect performance and optical modulation channel bandwidth. While there are many forms of complex modulation which offer improved utilization of the signaling channels bandwidth, solutions which offer effective backwards compatibility with previous signaling methods are preferred. Additionally, any change in symbol rate (which would emerge from choosing a different order of modulation at different media interfaces) would require a translator or gear-box to move between the dissimilar symbol rates. **Figure 21** illustrates several example paths forward from OIF's history with PAM4.

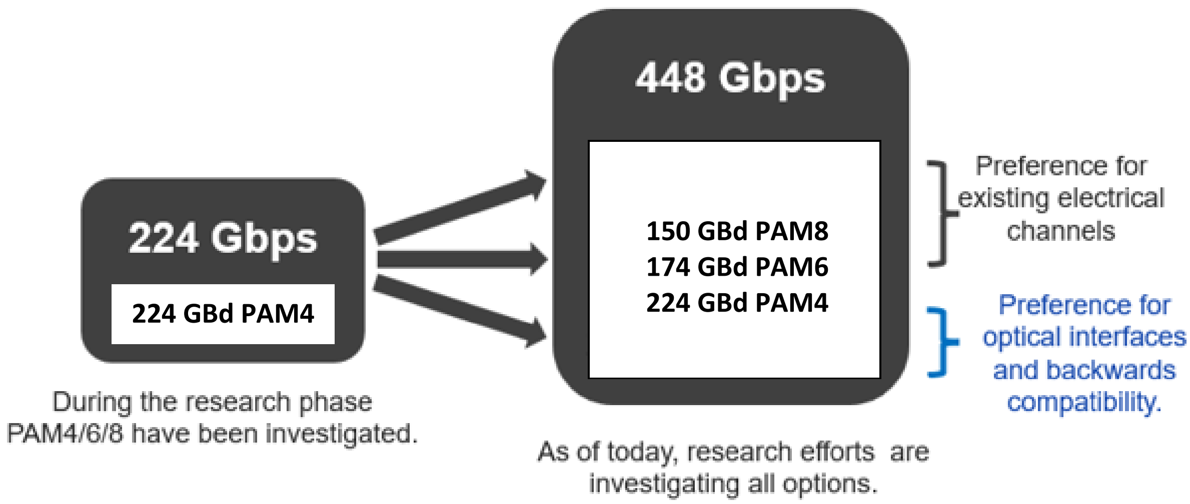

Figure 21 Modulation-order offers several paths to 448 Gbps

Higher-order PAM formats can be used in modern communication systems to increase data throughput without increasing the Nyquist frequency, which is beneficial for bandwidth-limited channels. However, as the number of levels increases, the distance between signal levels decreases, making the system more sensitive to noise. PAM8 requires significantly higher SNR than PAM4 to maintain the same BER, as illustrated in **Figure 22** and **Table 3**. Transmitter and receiver design complexity, and power consumption are also important aspects to consider when choosing the modulation format.

Figure 22 illustrates BER -vs- SNR curves of optimum receivers over an additive white Gaussian noise (AWGN) channel for three specific modulation formats under consideration for OIF's 448G project. The legacy 224G-LR and 224G-MR pre-FEC BER limit lines are illustrated on this chart for comparison purposes only. 448G will likely have different pre-FEC BER limit lines.

The ΔSNR [dB] relative to PAM4 is also illustrated in the three BER-vs-SNR plots of **Figure 22** and **Table 3**. It's clear that with the higher order of modulation there is a signal to noise penalty of 3.6 dB between PAM4 and PAM6, and a 6.1 dB SNR penalty for PAM8 relative to PAM4 at 10^{-5} BER.

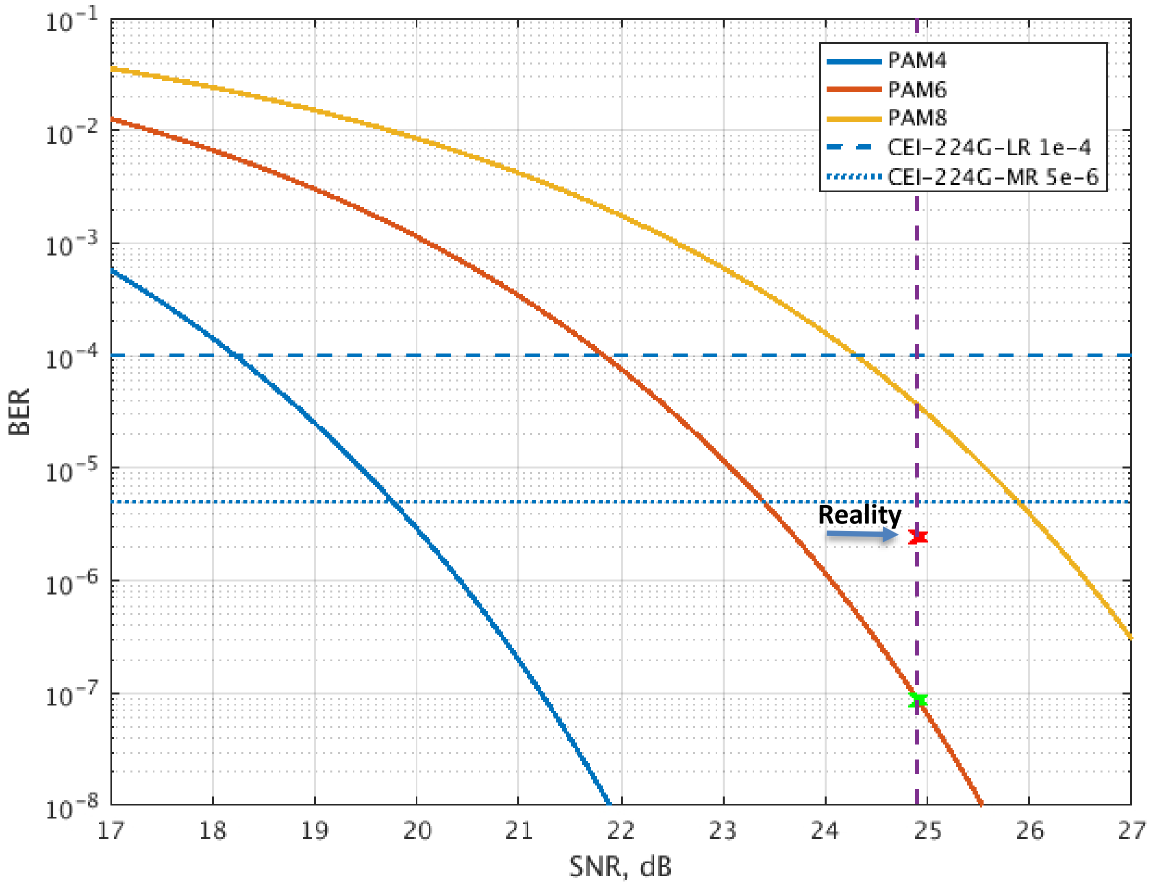

Figure 22 BER vs. SNR with optimum receivers over an additive white Gaussian noise channel

6.2 Physical measurements

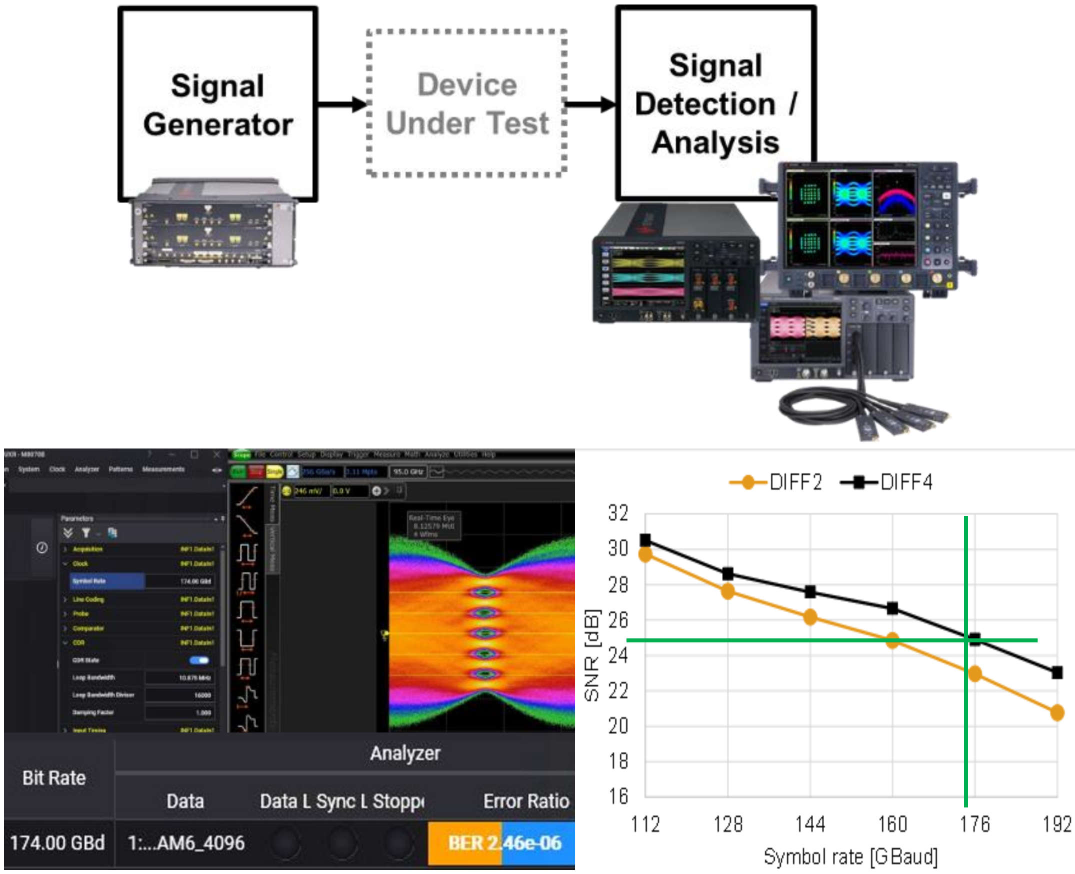

Figure 23 Comparison of physical measurements -vs- theory

Figure 23 illustrates an actual PAM6 instrument grade signal generation, channel propagation and acquisition. This test configuration offers a direct correlation to theory through several key measurements. The Signal Detection/Analysis system illustrated here is a real-time instrument receiving a PAM6 signal in this example. The pre-FEC BER at the receiver is 2.4e-6 BER illustrated as the red X on **Figure 22** which is well below most of OIF-CEI-224G pre-FEC limit lines. The PAM6 generator operating at 174 GBd (DIFF4 configuration in **Figure *24***) will deliver a signal with a SNR of 24.9 dB, as illustrated in **Figure *23***. Channel BW here is the theoretical limit (zero loss, skew, reflection etc.) which maps to 9E-8 BER and is the green X on **Figure 22**. The difference between 2.4E-6 and 9E-8 is over an order of magnitude which is significant, but this gap would be reduced with a proper reference receiver and receiver equalization. Everything in this example is un-equalized.

6.3 Signal generation

At this time, the most advanced signal generator available in a 256 Gbps Arbitrary Waveform Generator (AWG) can be interleaved or muxed in various configurations to achieve either an optical or electrical reference transmitter for PAM8, PAM6 or PAM4.

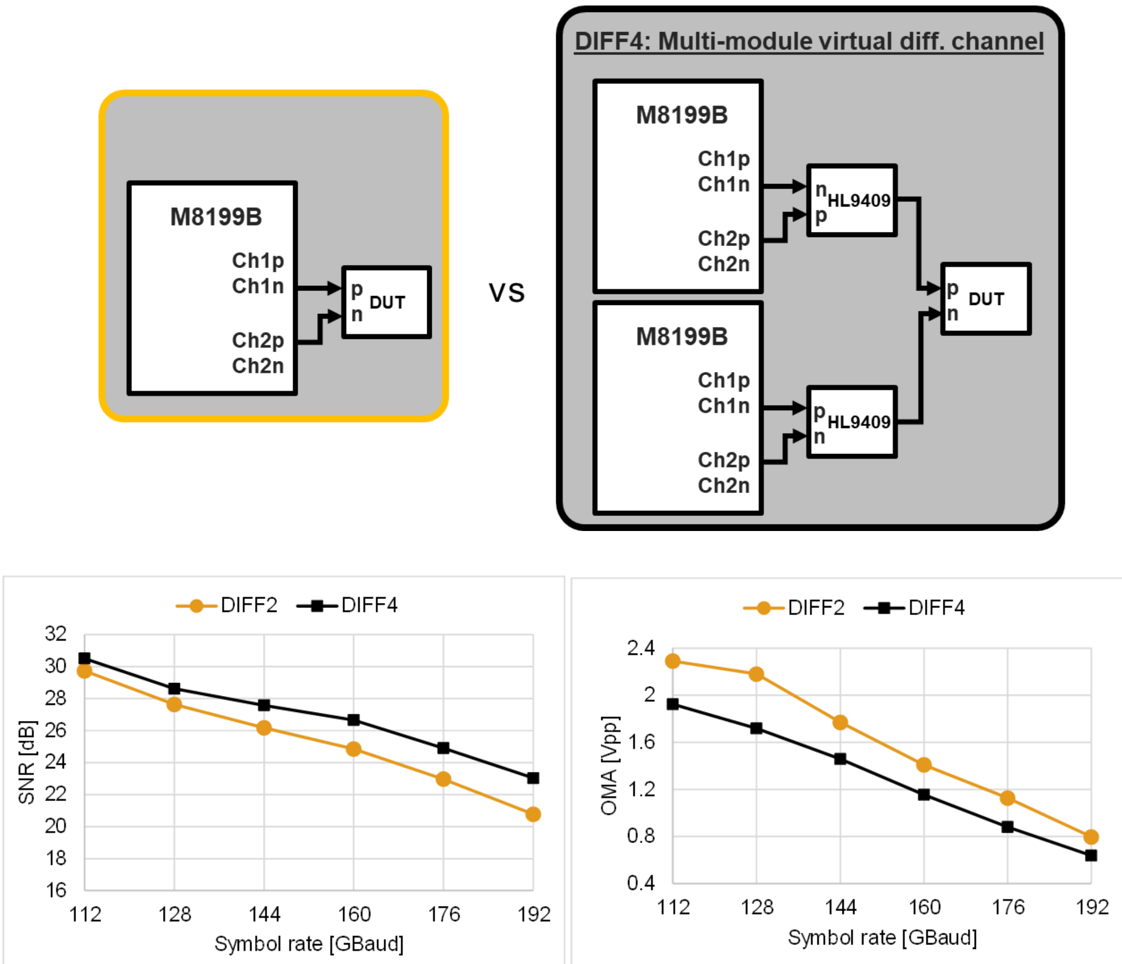

Figure 24 448G Generation through separate DACs to reject uncorrelated noise (improving SNR)

Figure 24 illustrates the incremental gains in SNR and OMA that can be achieved through interleaving adjacent DACs (Digital to Analog Converter) through a high bandwidth (>100 GHz) balun. This interleaving method provides variable signal D+/D- skew controls as well as serves as a key to reducing uncorrelated noise and improving SNR. Typical non compensated (no EQ) eye diagrams for this passive interleaver architecture are illustrated in **Figure 25**.

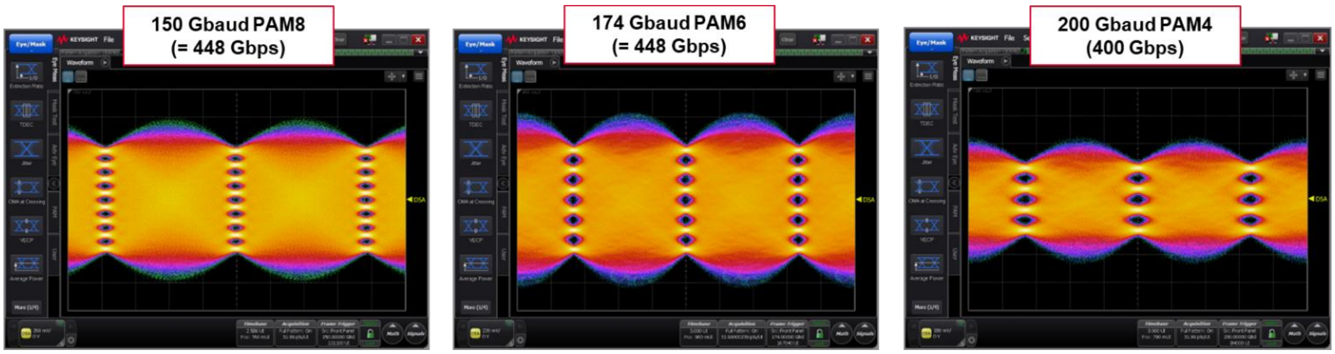

Figure 25 448G Differential signal generated with two M8199B AWGs (for each module, ch1p and ch2n connected to a high-bandwidth balun)

The operational limits of this passive interleaver system are evident in the SNR/Symbol rate plots of **Figure 24**. The M8199B AWG operates with a compensable bandwidth of approximately 81 GHz and its SNR and overall performance is limited to Symbol rates up to 192 GBd in its native configuration. Generation of a PAM4 signal at 224 GBd requires some heroics which we will discuss in the next section.

6.4 Signal generation heroics

Frequency Domain Interleaving methods using a banded upconverter and Heterodyne techniques allow us to bridge the gap between lower bandwidth generators and the need for speed associated with PAM4 signaling at 448 Gbps.

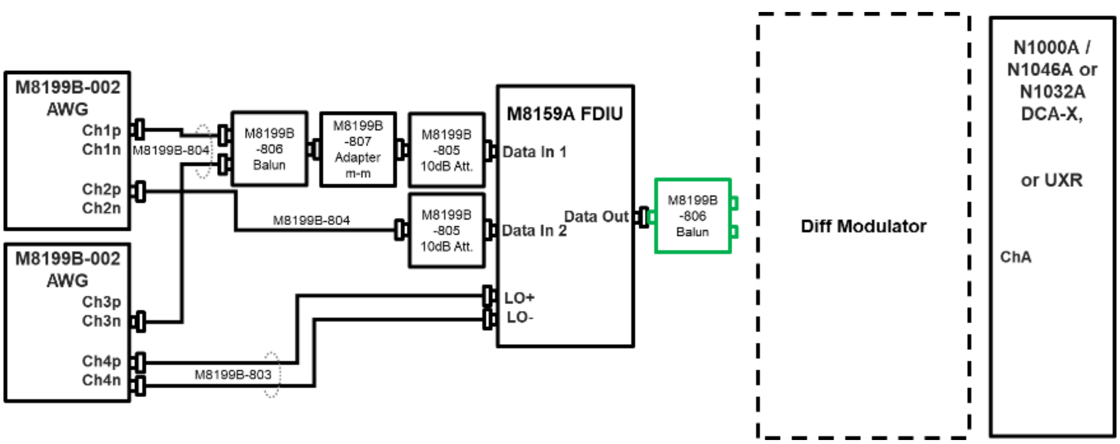

Figure 26 M8159A Frequency Domain Interleaver Unit (FDIU) setup

Figure 26 offers a pathway forward for research in higher speed modulation methods. Above 200 GBd PAM4 is where critical attention is required to compensable generator performance as well as the interconnect channel between the generator and the receiver. More on this later in section 6.5.

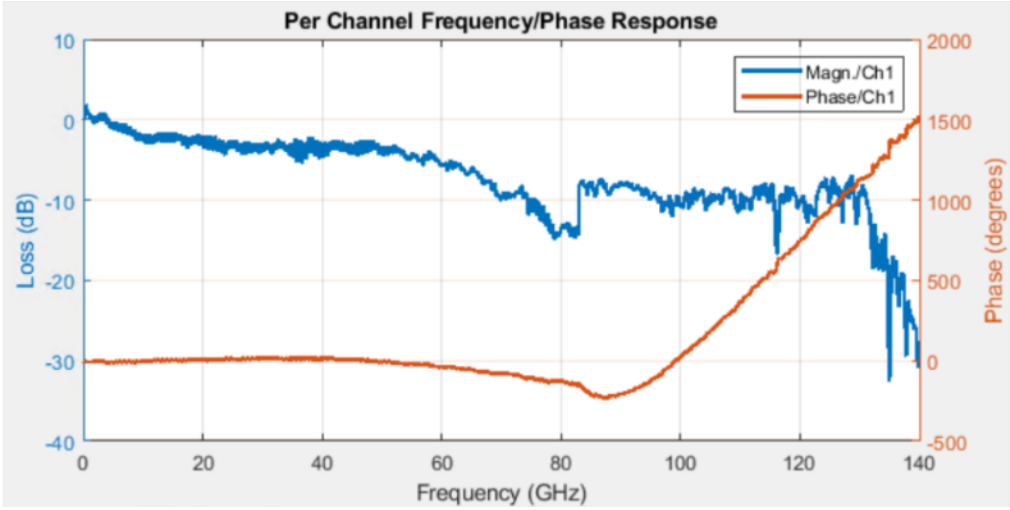

Figure 27 M8159A FDIU per channel frequency/phase response

Figure 27 illustrated the combined generator phase and magnitude response illustrating compensable performance up to 130 GHz. Signals at these speeds are pushing the performance limits of E/O modulators as well as acquisition system performances.

For the first time this configuration illustrates a clear path into the 200+ GBd PAM4 generation space with examples illustrated in **Figure 28** below.

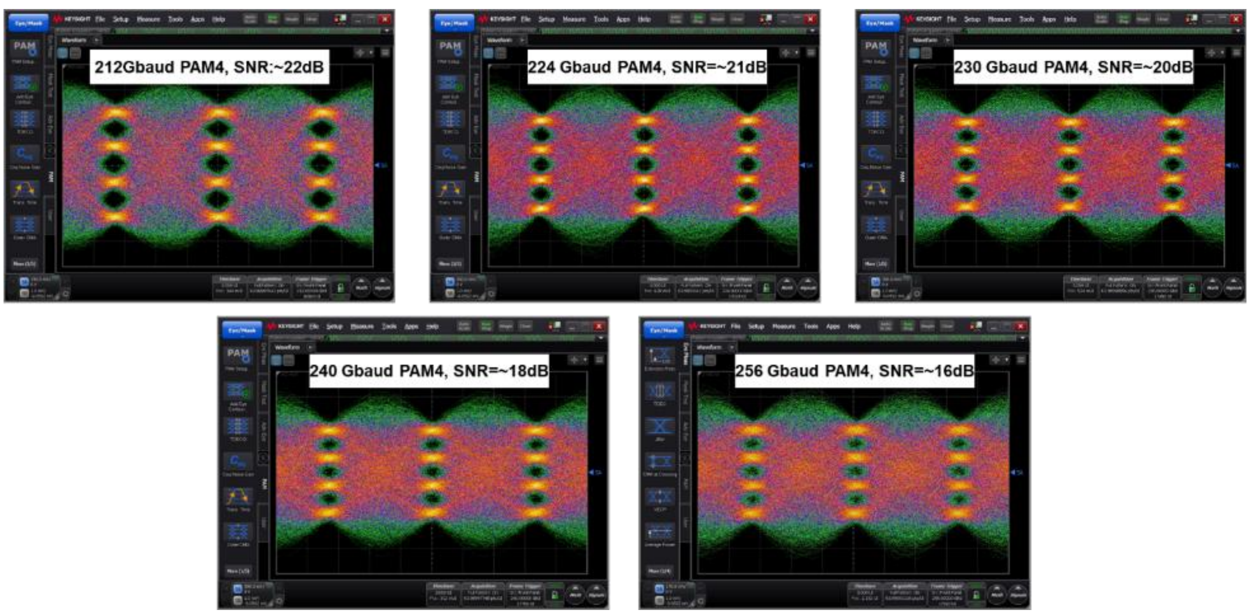

Figure 28 M8159A/M8199B > 200 GBd PAM4 examples

6.5 448G channel measurements and modulation considerations

The most compelling justification to move beyond PAM4 into a higher order of modulation is to reduce the signaling Nyquist frequency to something in line with contemporary mass produced high-performance data center (HPDC) interconnects. Today the state-of-the-art OSFP Gen2-Rev2 connector offers bandwidth up to 85 GHz. From **Table 3** we can see that this offers enough bandwidth headroom to support PAM8 and possibly PAM6 with no margin.

Figure 29 illustrates the HPDC state-of-the-art SDD21 (Insertion Loss) found in high density connector technology today.

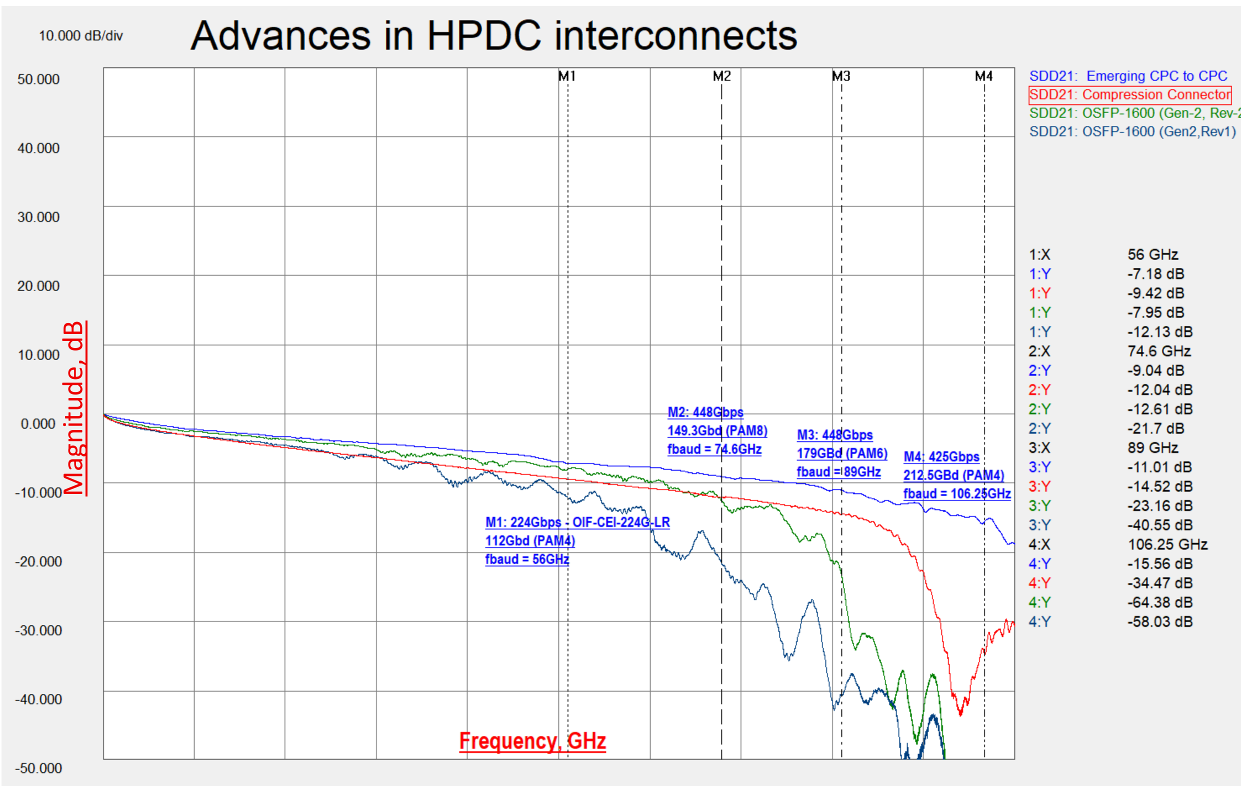

Figure 29 High Performance Data Center connector advances (SDD21 shown)

The green trace above is the differential insertion loss of an OSFP Gen2-Rev2 interconnect. Its operational limit is around 85 GHz, which is primarily driven by its return-loss exceeding -10 dB, beyond which the signal is largely non-compensable. The dark blue trace is an OSFP Gen2-Rev1 interconnect. Its operational limit was around 67 GHz and was what was the best available circa 2023. The red trace is representative of state-of-the-art compression connectors in late 2024, which extend the operational insertion and return loss by approximately 10 GHz to 100GHz, which is real progress, despite this not being a high volume HPDC connector.

The blue trace in **Figure 29** is the state-of-the-art (as of June 2025) 448G CPC interconnect performance, which for the first time is extending insertion and return loss numbers into operational levels beyond 106 GHz. This illustrates a clear avenue forward for CPC to CPC or substrate to substrate interconnects possibly supporting PAM4 at 224 GBd. Any CPC to modular interface will require at least a 25 GHz improvement to the current state-of-the-art OSFP to achieve 224 GBd PAM4 signaling at the modular interface.

6.6 448G measurement sciences summary

The velocity of advances in the physical layer measurement space has never been higher than it has been in 2025. Acquisition bandwidth, ENOB and Signal Generation capabilities will be central in supporting the decisions made by OIF-CEI regarding modulation formats and loss objectives for the emerging 448G IAs. Aside from the major decision around modulation, a few areas of focus are before us.

1. Advances in equalization methods. The expected channel loss numbers will likely require the publication of an MLSD component in the 448G reference receiver.
2. Co-Optimization methods in the compound reference receiver equalizer have resurrected the VEC or EECQ techniques as viable physical layer parametric properties that should likely be part of future specifications.
3. Crosstalk and associated bounded uncorrelated noise are a significant design parameter given possible edge speeds in the <3 ps region. This is a great reason to pursue the PAM6 or PAM8 route.
4. Connector technology is improving by 10 GHz roughly every 6 months. By mid-2026 PAM4 modulation may be possible through either new z-axis datacenter connectors or re-designed existing beam-on-blade designs today.
5. High performance instrumentation operating well above the current 120 GHz ceiling will be essential for effective analysis of 448 Gbps.

7 Opportunities for Future Work

This framework document identifies several areas where the OIF and other standards bodies could undertake future work to address the challenges and opportunities presented by 448G electrical interconnects. These include:

- Defining New CEI-448G IAs: The most direct follow-on work would be the development of new CEI-448G Implementation Agreements for specific application spaces (e.g., XSR, VSR, MR, and LR), specifying detailed electrical parameters, performance requirements, and test methodologies. This may involve exploring and standardizing new modulation schemes beyond PAM4, such as PAM6 or PAM8, if channel limitations persist.
- Channel Material and Connector Advancements: Collaboration with material scientists and connector manufacturers to drive innovation in PCB materials, package substrates, and connector designs to improve channel bandwidth and reduce insertion loss and reflections
- Advanced Equalization and DSP Techniques: Further investigation and standardization of advanced DSP-based equalization techniques, including noise and crosstalk cancellation, to achieve reliable performance over challenging 448G channels. This could involve defining new metrics or methodologies for evaluating the effectiveness of these techniques.
- Stronger FEC Schemes: Research and standardization of more robust Forward Error Correction (FEC) schemes for 448G electrical links, considering the trade-offs between coding gain, latency, power consumption, and implementation complexity. This may include evaluating new FEC architectures or enhancing existing ones.

8 Relation to Other Standards

These projects will potentially benefit from liaison activities with study groups and task forces and industry organizations including:

- ITU
- IEEE 802.3
- Fibre Channel
- InfiniBand
- JEDEC
- PCI-SIG
- CXL
- UEC
- UALink Consortium
- SNIA SFF
- Ethernet Alliance
- Open Compute Project (OCP)

9 Summary

Service providers, network customers and data center operators have clearly communicated that higher data rates like 448 Gbps are required to support scale-up, scale-out and front-end networks. These next generation data rates need to be implemented while also addressing challenges associated with power dissipation, density, performance, reach and cost. In addition, compatibility with legacy data rates and networks will be required in many applications. These goals can be achieved by having consensus amongst a broad cross section of component, subsystem, and system suppliers to leverage new technologies that drive signaling, architecture, and integration developments. As has been demonstrated in the past, most recently at 224G, the OIF is proposing to play a key role in coordinating industry activity to identify and develop critical technical solutions that will enable next generation data rates to be cost effectively deployed in the development of next generation equipment and networks.