WHITE PAPER

Performance Evaluation and Modulation Strategies for 448G Interconnects in Next-Generation AI Cluster Scaling

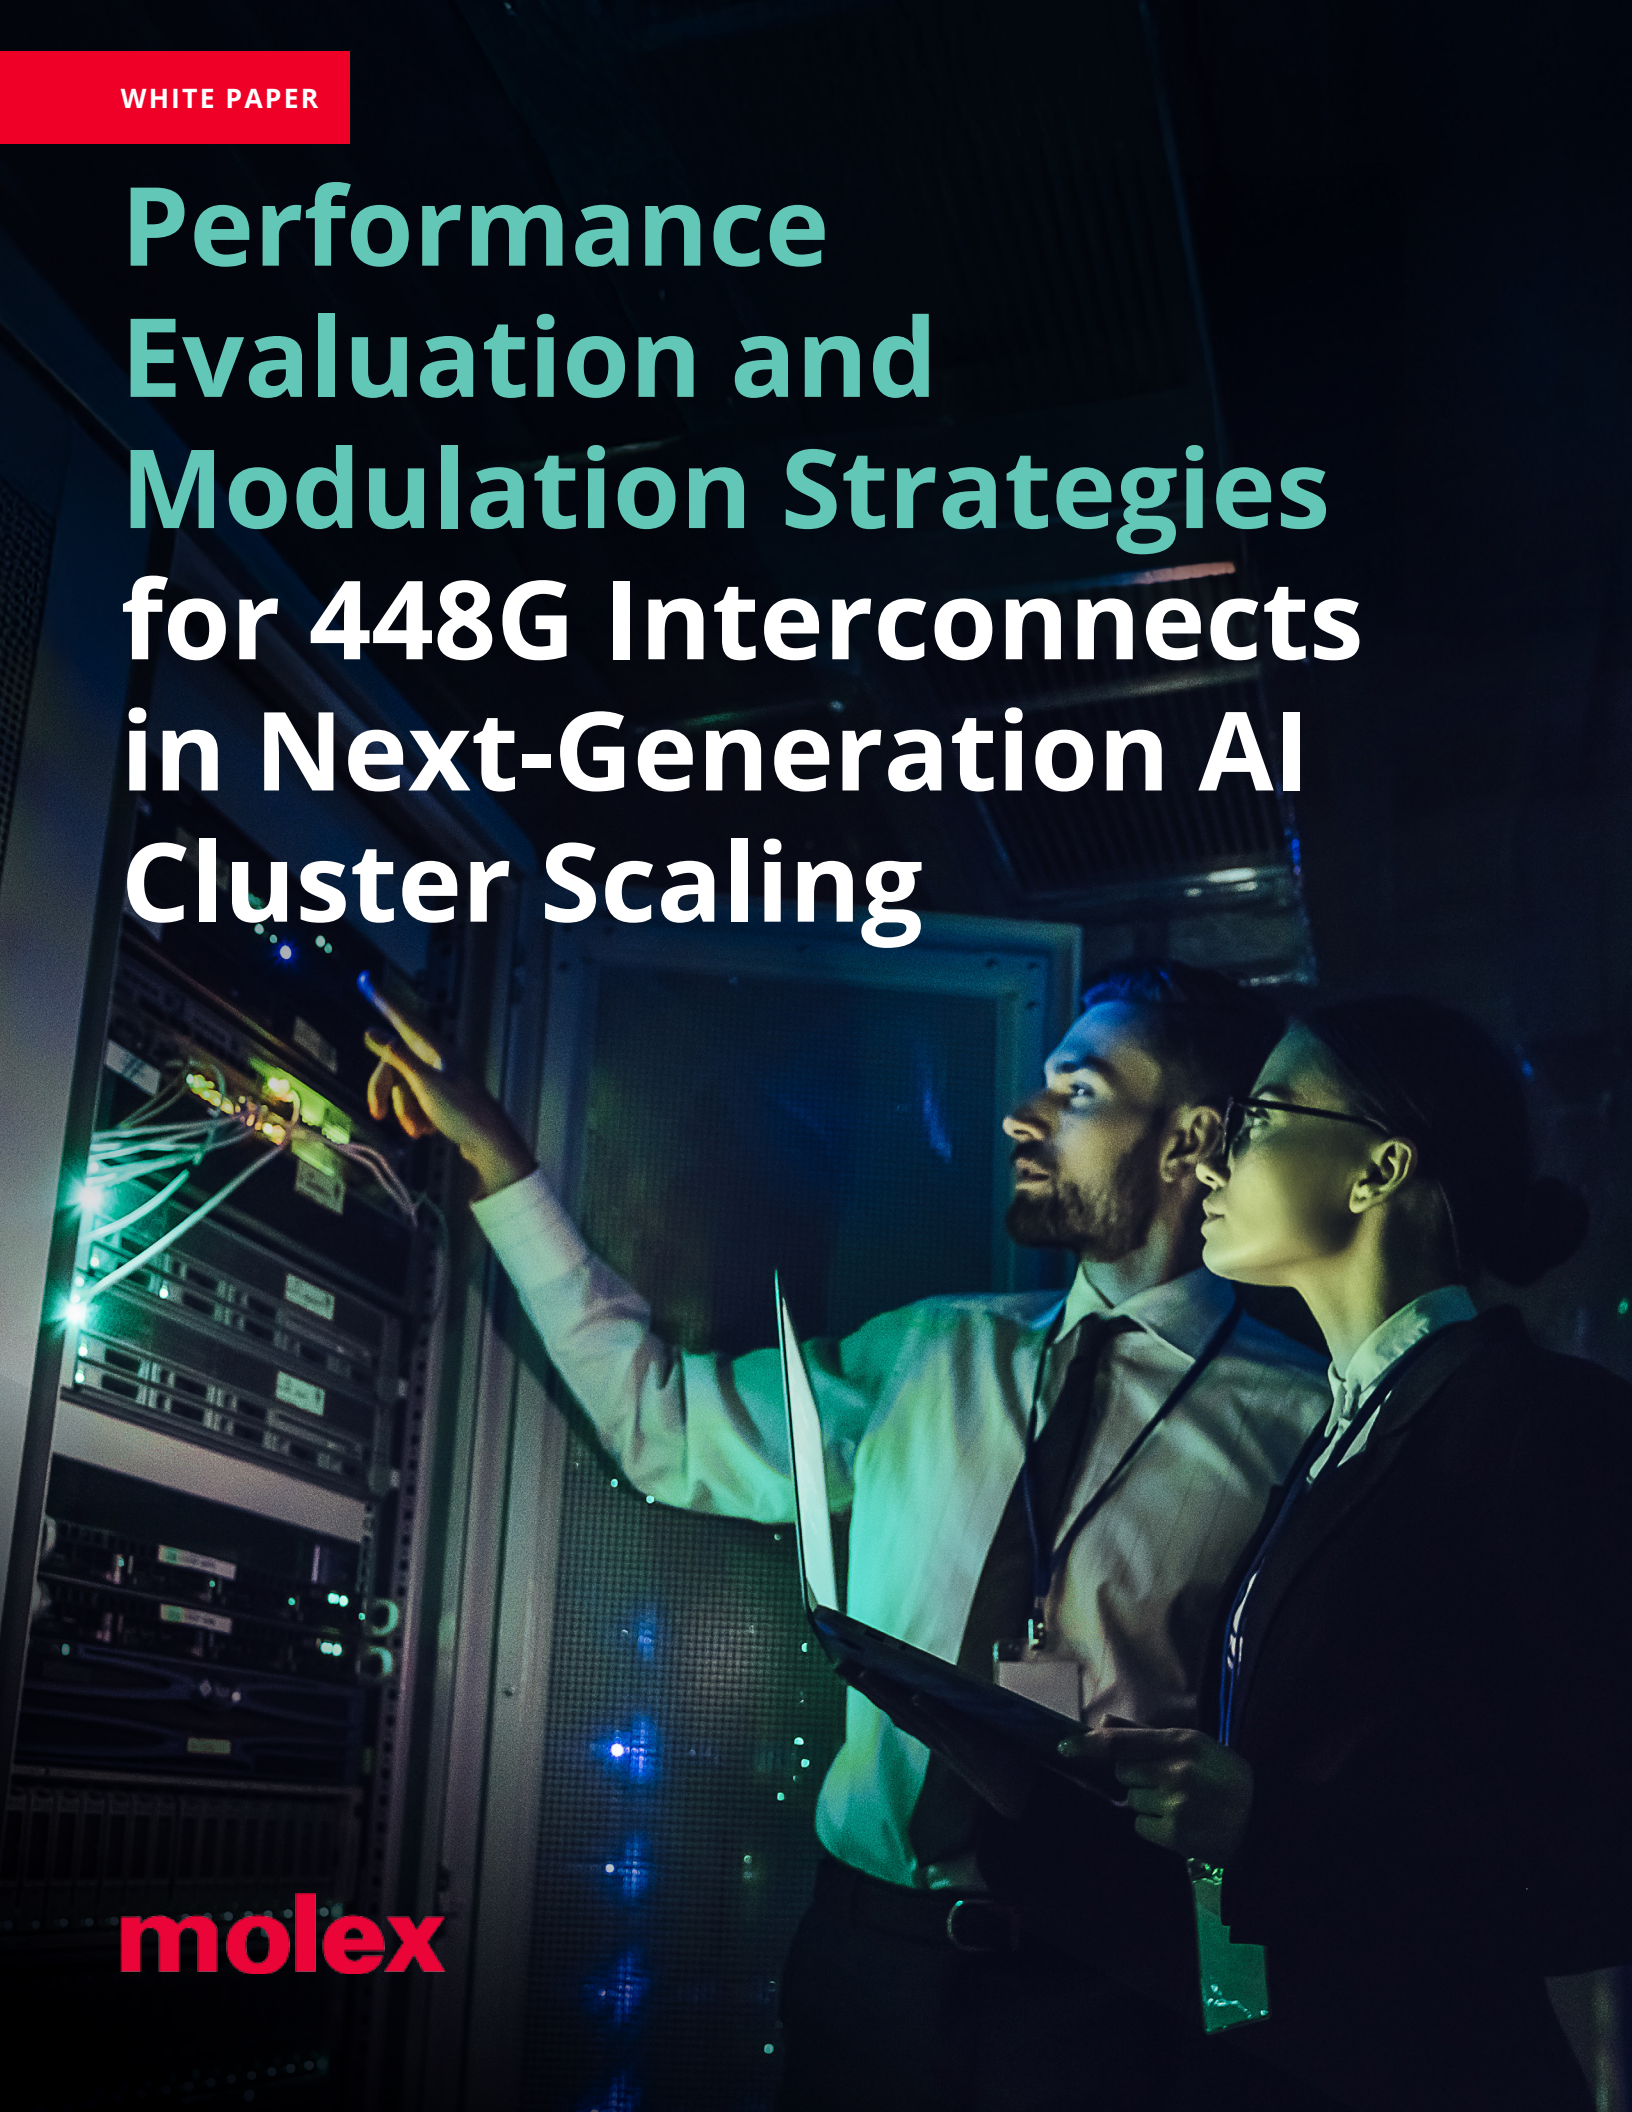

molex

AI is driving rapid investment in data center infrastructure, particularly as hyperscale operators expand the size and performance of AI clusters. The explosive growth in training workloads and inference applications has created new requirements for high-bandwidth interconnects to support thousands of interconnected compute nodes. Meeting these requirements depends on advancing serial link technology to 448G per lane, enabling future 3.2Tbps per port Ethernet interfaces.

Scaling AI clusters to this level places significant demands on the physical layer, where the interconnect fabric defines the achievable bandwidth density between processors, accelerators and network interfaces. Both optical and copper solutions are being considered, but copper interconnects remain attractive if design challenges in packaging, PCB layout and connector transitions can be overcome. Due to the bandwidth required to support candidate 448G modulation formats, it is still an open question whether traditional PCB interconnects, flyover interconnects and conventional connector form factors can be used at this data rate.

This paper presents a study examining copper interconnect performance at 448G using three proposed modulation schemes: PAM-4, PAM-6 and PAM-8. Each scheme carries different implications for loss-tolerance, linearity and equalization requirements, which in turn affect connector design and system architecture.

The study considers interconnect options between component packages and QSFP modules, providing an overview of signal integrity requirements that must be addressed to enable reliable 448G deployment in AI data centers.

SIGNAL INTEGRITY CHALLENGES AT 448G IN CHIP-TO-MODULE ARCHITECTURE

Scaling AI clusters places stringent requirements on the physical layer, where each interconnect must deliver higher bandwidth and minimal signal degradation. An earlier Molex simulation study comparing direct PCB routing to flyover twinax channels illustrates the limitations of conventional PCB routing with traces and vias. While direct PCB routing is simpler, it experiences significantly higher insertion loss over long reaches and elevated return loss at the PCB vias, resulting in pronounced insertion loss roll-off between 80 and 90 GHz. Flyover twinax connector designs exhibit similar insertion loss roll-off, but with much lower total insertion loss; thus, flyover cabling offers longer channel reach.

Co-packaged copper (CPC) has attracted considerable interest as a solution to circumvent the high losses incurred during signal transitions through PCB and connector vias. These modules connect to trace routing inside the package substrate and interface directly to twinax cabling, bypassing the bandwidth-limiting PCB vias. This architecture, called chip-to-module (C2M), is outlined in Figure 1.

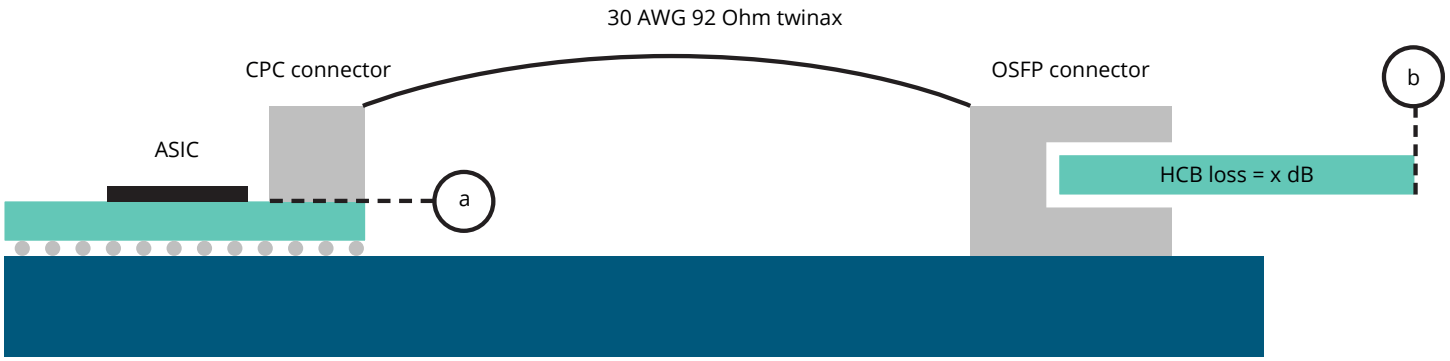

Figure 1: C2M architecture using co-packaged copper connectors with flyover cabling connecting to an MSA-standardized connector and module

Connectors for co-packaged copper also present signal integrity challenges related to the surface mount device (SMD) land pattern, contact stub lengths and routing through vias from the package into the connector. However, once the signal traverses the connector and passes along the cabling, much lower insertion loss is achieved and coaxial connector designs (e.g., twinax) tend to have much lower return loss than PCB interconnects. This applies to flyover cabling to QSFP/OSFP connectors and to copper backplane connectors.

The signal integrity performance metrics depend on the modulation format, as different formats require specific power margins between signal levels in the serial bitstream. Table 1 and Table 2 compare the channel bandwidth requirement against the channel reach, bit error rate (BER) and signal-to-noise ratio (SNR) for each proposed modulation format. As the current technology being deployed in AI data centers is 224Gbps-PAM-4 with 56 GHz channel bandwidth, the relative values are compared against 448Gbps-PAM-4 with a 112 GHz required channel bandwidth.

Modulation	Dimensions	Bits per Dimension	Signaling Rate, GBd	Bandwidth, GHz	Distance Reduction, dB
PAM-4	1	2	225 (212.5)	112.5 (106.25)	—
PAM-6	2	2.5	180 (170)	90 (85)	-4.44
PAM-8	1	3	150 (142.5)	75 (71.25)	-7.36

Reduced bandwidth requirements

Increased difficulty to detect signals

Table 1: Signaling characteristics and channel reach for the candidate modulation formats in 448G per lane interconnects

Modulation	BER at SNR = 19 dB	Δ	SNR for BER = 2.4e−5	Δ
PAM-4	2.4e−5	—	19	—
PAM-6	3e−3	125x	22.6	+3.6
PAM-8	1.5e−2	625x	25.1	+6.1

Table 2: BER and SNR requirements for each modulation format shown in Table 1

Table 2 provides a comparison of BER required given the target SNR value in PAM-4 modulation, and vice versa. The data shows that multiple factors influence the SNR observed at the receiver side of an interconnect:

- Insertion loss
- Return loss (reflections)
- Crosstalk
- Skew/jitter

The BER and SNR requirements for each modulation format in the higher bandwidth channels illustrate how, with N-ary PAM, lower bandwidth channels also require lower reflections and crosstalk for a given BER target due to the higher constellation density. Improved detection, link equalization or error correction can be used to relax SNR requirements.

At these data rates, mechanical features in the connector can limit overall channel performance. Factors like stub length, solder geometry, package-to-connector transitions and connector-to-twinax transitions can create reflections at low frequencies and resonances at high frequencies. Collectively, these create a roll-off in the insertion loss at high frequencies, effectively defining the upper limit of the channel bandwidth. High pin density on the connector modules also creates a risk of crosstalk at the package-to-connector transitions. In addition to the insertion loss, crosstalk will determine the SNR and BER values at the receiver side of the interconnect.

STUDY ON THE USE OF CO-PACKAGED COPPER FOR 448G INTERCONNECTS

The signal integrity characteristics observed in the previous simulation study motivate investigation of co-packaged copper for 448G interconnects at various modulation frequencies. The clear relationship between modulation format, losses, and SNR and BER limits warrants an experimental study involving real channels in a typical deployment.

The remainder of this paper presents an analysis of signal integrity metrics in 448G channels employing a standard link architecture. The study aims to achieve the following objectives:

- Evaluate SNR, BER and insertion loss for PAM-6 and PAM-8 signaling
- Determine bandwidth limits based on rollover in the insertion loss
- Examine crosstalk within the identified bandwidth range for interconnects in the C2M architecture
- Compare signal propagation in the C2M architecture in the host-to-module and module-to-host directions

Real co-packaged copper and OSFP connectors were used in the C2M architecture. Insertion loss, BER, and SNR margin were measured and compared for PAM-6 and PAM-8 signaling.

The link architecture examined in this study and the estimated BER limits in each section of the architecture are shown in Figure 2. In this C2M architecture, very short reach (VSR) channels of 300mm and 500mm length were examined in both host-to-module and module-to-host directions, yielding results on crosstalk and insertion loss along an interconnect. Then, the channel reach and modulation format were varied while the SNR and BER were measured in the presence of crosstalk. These results provide sufficient data to compare the performance of channels supporting PAM-6 and PAM-8 modulation.

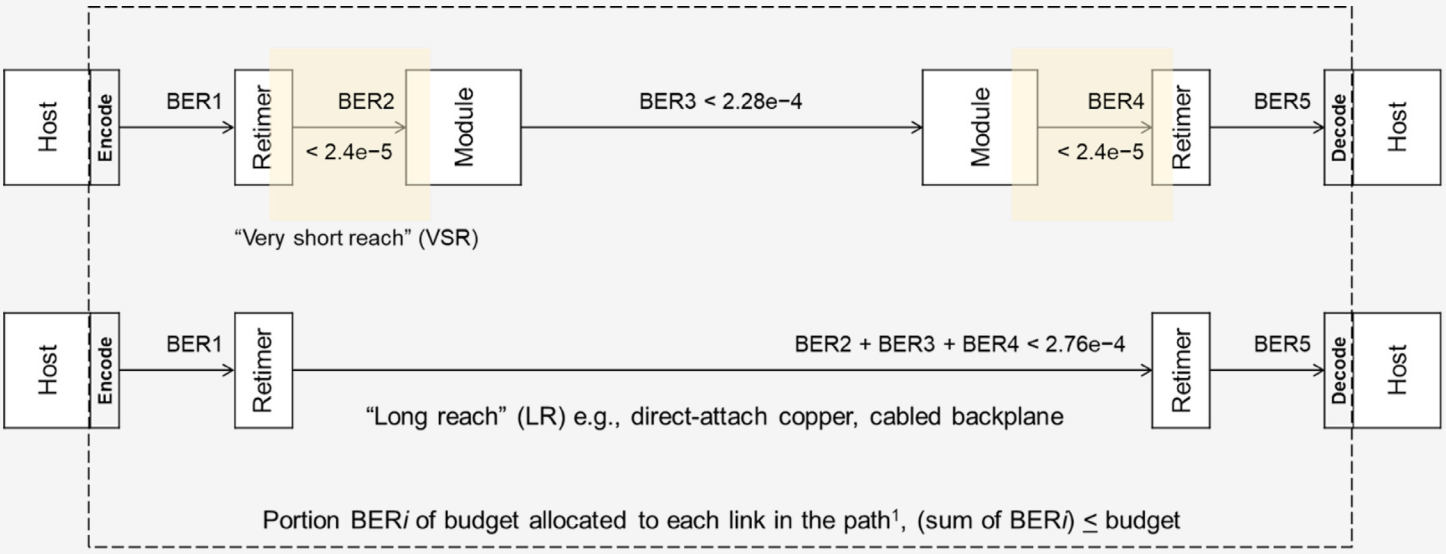

Figure 2: Link architecture investigated in the current study

[1] Link BER may need to be less than BER*i* if errors occur in a way that impairs the performance of the decoder.

STUDY FINDINGS AND CONCLUSIONS

Insertion Loss and Crosstalk

Figure 3 shows the insertion loss and crosstalk curves for host-to-module and module-to-host signal propagation. Both curves show strong insertion loss roll-off, pinning the upper channel bandwidth limit near 90 GHz. This confirms that these channels can support PAM-6 and PAM-8 signaling at 500mm channel reach.

The results pertain to channel TX6 (host-to-module) and RX6 (module-to-host) on a standard OSFP transceiver module, specifically chosen for exhibiting worst-case crosstalk conditions. Crosstalk levels remain low across the channel bandwidth, with PSNEXT in the host-to-module direction only rising above -60 dB at frequencies approaching 80 GHz. PSFEXT in the host-to-module direction only occasionally rises above -50 dB starting at around 75 GHz.

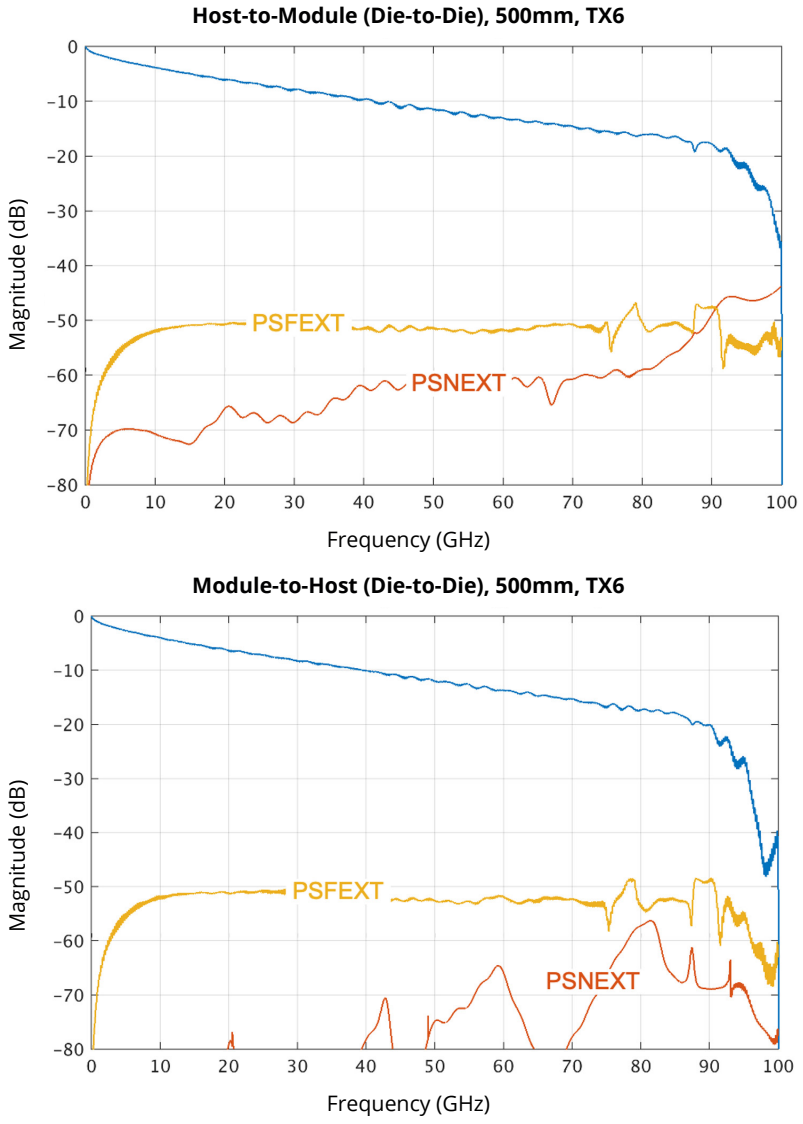

Figure 3: Die-to-die insertion loss and PSFEXT/PSNEXT results in the host-to-module (top panel) and module-to-host (bottom panel) direction; channel length = 500mm

BER and SNR Margin

Given the potential crosstalk levels in these channels, as shown in Figure 3, it is critical to examine expected BER and SNR margin. Table 3 and Table 4 show the insertion loss at Nyquist frequency, observed BER and available SNR margin for PAM-6 and PAM-8 formats. Channels TX6/TX8 and RX6/RX7 were chosen for this evaluation because they present the worst-case crosstalk.

Host-to-Module		PAM-6				PAM-8			
Pair	Cable Length, mm	Signaling Rate, GBd	Insertion Loss, dB	BER	SNR Margin, dB	Signaling Rate, GBd	Insertion Loss, dB	BER	SNR Margin, dB
TX6*	300	170	15.1	4.8e−7	1.7	145	13.7	1.8e−5	0.2
TX8**	300	170	15.1	4.4e−7	1.7	145	13.7	1.7e−5	0.2
TX6	500	170	16.5	5.4e−7	1.6	145	14.9	2.4e−5	0
TX8	500	170	16.5	4.9e−7	1.7	145	14.9	2.1e−5	0.1

*Channel with worst-case FEXT
**Channel with worst-case NEXT

Table 3: Host-to-module insertion loss, BER and SNR results

Module-to-Host		PAM-6				PAM-8			
Pair	Cable Length, mm	Signaling Rate, GBd	Insertion Loss, dB	BER	SNR Margin, dB	Signaling Rate, GBd	Insertion Loss, dB	BER	SNR Margin, dB
RX6*	300	170	16.5	3.9e−7	1.8	145	14.8	1.8e−5	0.1
RX7**	300	170	16.9	4.2e−7	1.7	145	14.9	1.7e−5	0.2
RX6	500	170	17.9	5.4e−7	1.6	145	15.9	2.3e−5	0
RX7	500	170	18.3	5.3e−7	1.6	145	16	2.1e−5	0.1

*Channel with worst-case FEXT
**Channel with worst-case NEXT

Table 4: Module-to-host insertion loss, BER and SNR results

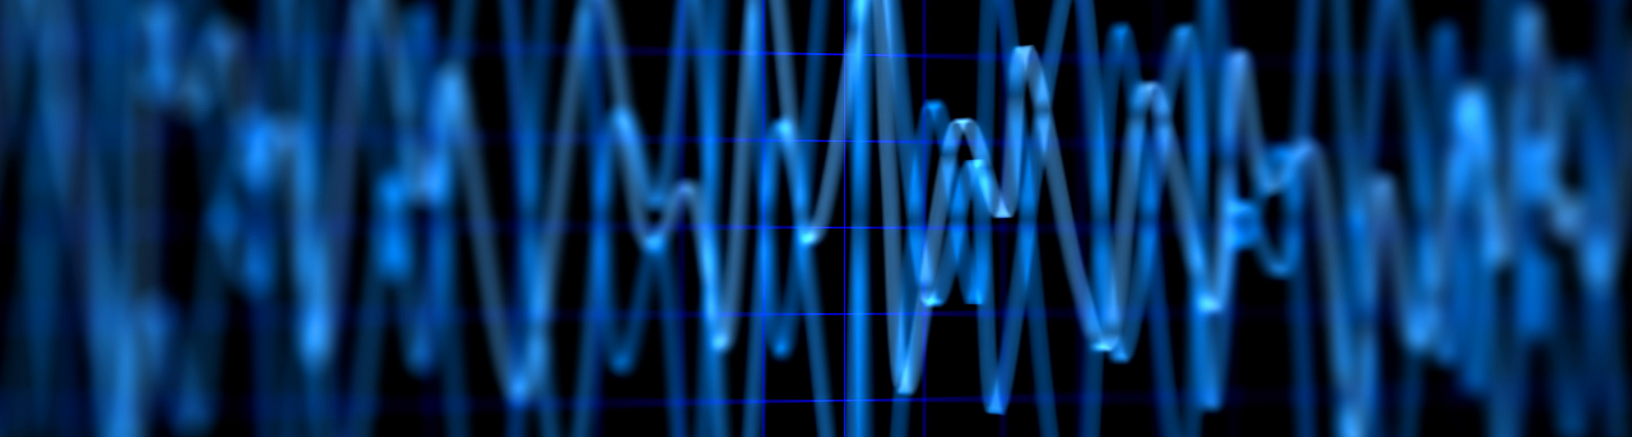

The results demonstrate that both PAM-6 and PAM-8 are technically feasible for use in these channels with OSFP connectors at 300mm and 500mm lengths. Because PAM-8 has the expected higher BER, this leaves almost no remaining SNR margin relative to the upper BER limit of 2.4e−5 at the retimer/module interface (see Figure 2). This leaves the link vulnerable to additional sources of noise, such as crosstalk, external electromagnetic interference (EMI) or power rail fluctuations.

448Gbps-PAM-4 modulation is not feasible using traditional OSFP paddle card interfaces due to insertion loss roll-off below the required 112 GHz channel bandwidth. Using PAM-4 at this data rate would require multi-source agreement (MSA) standard changes to increase the roll-off threshold.

Based on these results, PAM-6 emerges as the most viable option, offering a better balance of BER, SNR and bandwidth margin. The required channel bandwidth coincides with the insertion loss roll-off observed in the C2M architecture, which motivates some changes to the MSA standard for OSFP connectors and paddle card modules.

Proposed Changes to the MSA Standard

Based on the above results, several updates to the paddle card interface could help further enable PAM-6 at 448G-per-lane data rates. These involve connector and pad layout changes on the fiber module PCB to extend channel bandwidth in the C2M architecture beyond 90 GHz. The proposed changes are summarized in Table 5.

Proposal	Detail
Add chamfer to the module card edge to facilitate shortened signal beam tips	Increase chamfer from 0.25mm to 0.30mm
Reduce the module card signal pad length	Increase the nominal distance from card edge to signal pad edge from 1.70mm to 1.90mm
Reduce the wipe length tolerance to facilitate shortened module card signal pads	Tighten from +/-0.395mm to +/-0.200mm
Require additional exposed length of ground pads on the module card	Increase minimum ground pad length from 1.40mm to 2.50mm
Modify the OSFP connector pad layout on the host PCB	Use Ø0.36mm with via-in-pad for differential signals
Eliminate the pad array offset on the top and bottom sides of the fiber module PCB	Align the pad layout on top and bottom of module cards

Table 5: Proposed changes to the MSA standard for paddle card interfaces to support 448Gbps-PAM-4 and higher order N-ary PAM

PAVING THE PATH TO 448G

This study demonstrates that traditional paddle card interfaces with OSFP connectors and co-packaged copper connectors in a C2M architecture can support 448G signaling using PAM-6 or PAM-8 modulation. The BER and SNR results show PAM-6 as the preferred modulation format, contingent upon MSA updates to paddle card and connector designs. Without these changes, PAM-8 may remain feasible; however, enhancements such as inner coding or advanced equalization will be necessary to improve the SNR margin.

Molex is paving the path to 448G interconnect technology through extensive research and deep engineering expertise, building on our proven leadership at 112G and 224G speeds. By advancing connector architecture and signal integrity, Molex empowers data centers to move information faster and with greater signal clarity, meeting the performance demands of emerging AI-powered data environments.

To learn more about the foundational technologies enabling this transition, including design strategies for next-generation data centers, visit our 224Gbps-PAM-4 High Speed Data Center Technology Page.

Order No. 987652-9471 USA/LX/2025.11