SYNOPSYS®

WHITE PAPER

Designing for 448G: Modulation, DSP, and Channel Trade-offs in High-Speed SerDes

Authors

Kent Lusted
Distinguished Architect, Synopsys

Priyank Shukla
Director of Product Management, Synopsys

Disclaimer:

"At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position of IEEE."

Introduction

As the industry moves toward deploying 400 Gbps/lane Ethernet and 448G SerDes, system designers face new architectural challenges across host interface architectures, channel design, modulation selection, and DSP performance scaling. This whitepaper focuses on the electrical and implementation-level feasibility of 448G signaling in the context of AI and HPC cluster networks. It examines how host architectures—retimed and unretimed—impact signal integrity and power, explores the trade-offs between PAM4 and PAM6 modulation, and presents detailed channel simulations and DSP design implications using real-world 448G electrical topologies. It also offers insight into equalization strategies, ADC/DAC resolution needs, and MLSD scaling, supported by recent contributions to the industry from Synopsys.

To understand the broader system-level drivers, standardization efforts, and how scale-up and scale-out AI architectures are shaping the need for 448G, read our companion paper: "Enabling the 448G Era: System Architecture and Standards for Next-Gen AI Networks." That paper provides the full context for why 448G is emerging now—and how ecosystem alignment is accelerating deployment timelines.

What will be the Influence of Host System Architecture at 448G?

The transition to 448G signaling marks an important milestone for Ethernet-based networks. Unlike previous generations where the evolution of Ethernet was primarily driven by front end networks, 448G introduces the industry's first significant deployment of accelerator-to-accelerator (XPU-to-XPU) links within AI-driven clusters. As a result, host system architecture becomes a crucial factor influencing channel characteristics, trade-offs, and the overall design choices for PHY implementations.

At the host level—the System-on-Chip (SoC) or ASIC initiating the packet transmission—two primary architectures are being evaluated by system implementers: unretimed and retimed interfaces. Each approach presents unique advantages and trade-offs regarding power consumption, complexity, and signal integrity.

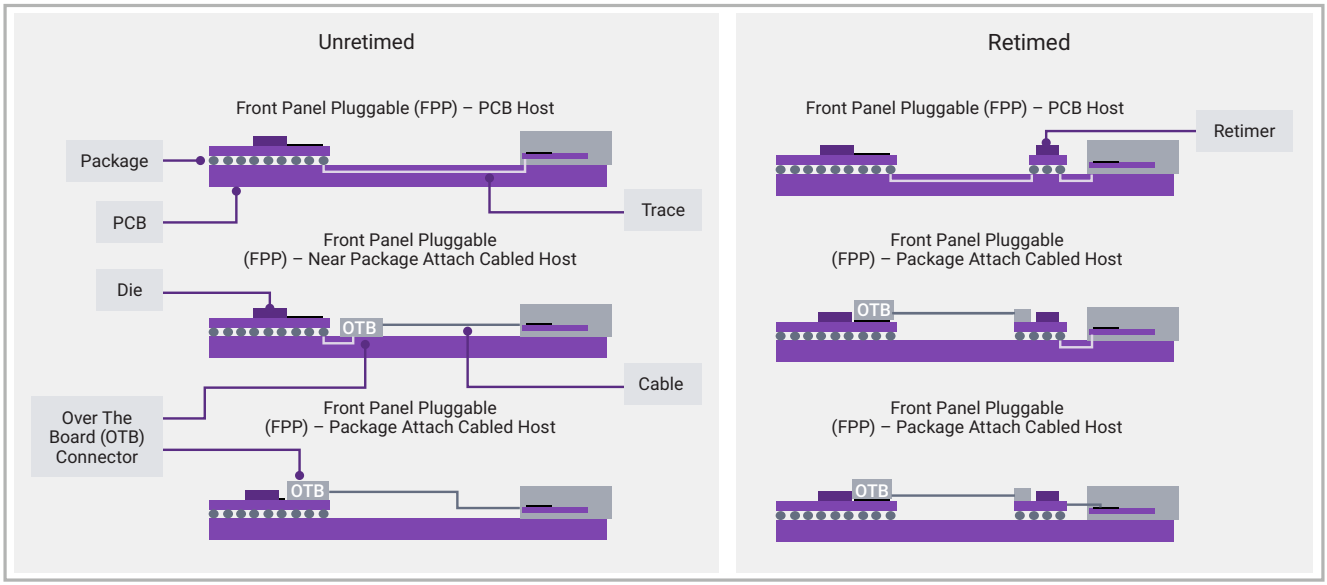


Figure 1: Host System Architectures - Unretimed vs. Retimed

Unretimed Host Interfaces

Unretimed host system architectures involve direct electrical connections from the host device to the endpoint without intermediate retiming components. By omitting retimers, these designs significantly reduce power consumption, making them highly attractive for power-sensitive applications.

There are several unretimed configurations currently under evaluation:

- **Conventional PCB Host (Front-Panel Pluggable, PCB-based):** The SoC is mounted directly onto a printed circuit board (PCB), with electrical traces routing signals directly to a front-panel pluggable connector. This conventional method offers simplicity but becomes increasingly challenging at higher data rates due to signal integrity issues
- **Near-Package Attached Copper Cable Host:** As demonstrated at 224G, copper cables attached closer to the silicon package help minimize crosstalk and reflections, reducing the signal processing burden on SerDes. This architecture is well-positioned for adoption at 448G, offering improved channel performance with moderate complexity
- **Direct Front-Panel-to-SoC Cabled Host:** Direct cable connections from the front panel connector to the SoC further shorten the electrical path, potentially extending reach capabilities and simplifying SerDes equalization requirements at higher signaling rates

Retimed Host Interfaces

Retimed host system architectures incorporate retimers in the signal path between the host SoC and the endpoint. While adding retimers increases overall power consumption and complexity, it significantly improves signal integrity, extends channel reach, and eases design constraints on the host system and the SerDes. This makes retimed architectures suitable for longer reach scenarios or highly demanding signal integrity conditions.

Implementers evaluating retimed architectures must carefully balance the benefits of extended reach and improved signal quality against increased power consumption and design complexity.

Ultimately, the choice between retimed and unretimed architectures at 448G will depend heavily on specific deployment scenarios, power budgets, and performance requirements of targeted AI-driven clusters.

Modulation Choices and Trade-offs for 448 Gbps Signaling

Choosing an appropriate modulation scheme is a critical decision in deploying 400 Gbps/lane Ethernet, influencing key parameters such as signal integrity, reach, backward compatibility, power consumption, and overall PHY complexity. From a SerDes design perspective, there are numerous modulation options currently to be consideration by the industry, each with distinct advantages and challenges.

Table 1. Comparison of Modulation Schemes

Modulation	Levels	Constellation	Baud-rate	Nyquist-rate	Consideration
PAM4	4	1D	224Gbd	112GHz	Backwards compatible, aligned with optics
PAM6	6	1D	179.2Gbd	89.6GHz	Offers slight bandwidth relief, has SNR penalty
CROSS-32	6	2D	179.2Gbd	89.6GHz	CROSS-32 has slightly more detector margin than PAM6 for 50% of received symbols, 6-level constellation, more complex detector (2D)/MLSD
PR-PAM4	7	1D	224Gbd	~56GHz	Reduced noise immunity on link, reduced SNR due to more levels, may need training sequence for adaptation
PAM8	8	1D	149.3Gbd	74.6GHz	High SNR penalty, more relief on BW
DSQ-32	8	Double Square 2D	179.2Gbd	89.6GHz	No different than cross-32 in detector margin, 8-level constellation vs 6-level
BiDi-PAM4	4	1D x 2	112Gbd	56GHz	Bi-directional differential links Need Hybrid/echo canceller (but keeps same I/O count)
SE-PAM4	4	1D x 2	112Gbd	56GHz	Single-ended links, need lane to lane alignment Sensitive to common mode and xtalk (xtalk canceller) (but keeps same I/O count)
DMT	PAR	QAM	~224Gbd	~112GHz	Need FFT/inv-FFT, not backwards compatible

PAM4: Proven Compatibility and Alignment with Optics

PAM4 modulation, widely adopted at 100 Gbps/lane and 200 Gbps/lane Ethernet speeds, remains a highly attractive option for 400 Gbps/lane Ethernet. The industry's familiarity with PAM4 ensures extensive backward compatibility, easing the transition from current deployments. PAM4 could also align naturally with optical modules, such as 400 Gbps/wavelength optics, simplifying implementations of electrical-optical conversions.

However, achieving desired reach and signal integrity at 448G signaling rates may challenge channel designs, potentially increasing SerDes complexity and power requirements.

PAM6: Lower Nyquist Frequency and Extended Reach

PAM6 modulation, another single-dimensional amplitude modulation scheme, offers the advantage of a lower Nyquist frequency (approximately 90 GHz at 448G signaling). This relaxation in bandwidth requirements significantly eases channel design constraints, allowing longer channel lengths or lower-quality channels to achieve target performance.

The trade-off, however, is increased complexity in digital signal processing (DSP), higher power consumption, and more intricate silicon implementations. Despite these challenges, PAM6 is receiving substantial industry attention, particularly for purely electrical networks aiming for extended reach without resorting to optical interconnects.

Alternative Modulation Schemes: Bidirectional Modulation

Other modulation approaches, such as simultaneous bidirectional (SBD) signaling, are under investigation due to their potential for efficient bandwidth utilization and reduced lane count. However, these advanced techniques currently lack extensive silicon validation, making first-pass silicon success more challenging. As a result, they remain more suitable for proprietary interconnects or longer-term exploratory research rather than immediate open-ecosystem deployment.

Considering open standards and widespread industry adoption, PAM4 and PAM6 modulation schemes currently represent the most viable and widely supported options. Most ecosystem consensus is observed around these two schemes, with PAM4 strongly preferred for backward compatibility and optical alignment, and PAM6 emerging as a compelling option for networks prioritizing extended electrical reach and relaxed channel constraints.

Ultimately, the choice among these modulation techniques depends heavily on the system-level trade-offs that implementers are prepared to make regarding reach, channel complexity, power, and silicon maturity.

Synopsys Contributions to 448G PHY Ecosystem

As a leading Silicon IP provider with deep SerDes expertise and an ecosystem enabler, Synopsys is actively contributing to the industry's transition to 400 Gbps/lane Ethernet. Recognizing the need for detailed PHY and system-level studies, Synopsys has initiated early analysis and simulations using available 224G channel models to assess their viability at 448G signaling rates.

In a significant contribution made in January 2025, Synopsys demonstrated system margin feasibility for both PAM4 and PAM6 modulation schemes for existing 224G channels. These studies illustrated potential trade-offs in channel performance and equalization complexity, highlighting the viability of dual-simplex PAM4 and PAM6 as starting points for detailed industry considerations.

Since then, additional contributions from channel vendors have expanded available models for 448G, allowing Synopsys to refine simulations and support further detailed analyses.

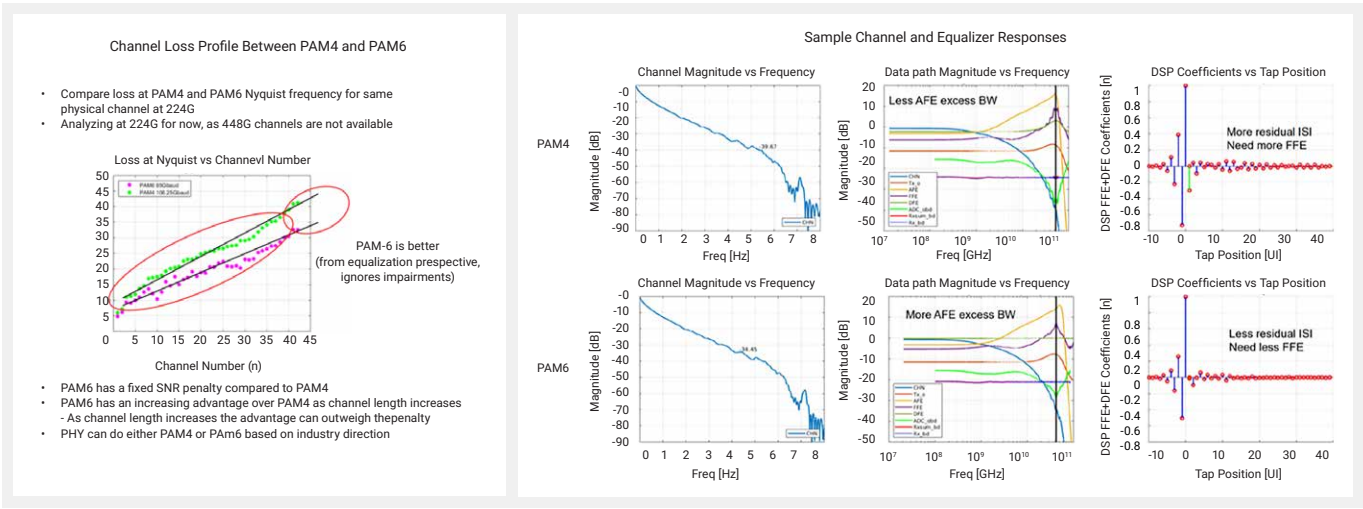


Figure 2: DSP Considerations for 400 Gbps by Kent Lusted et al, Synopsys at SNIA_SFF 400G "AI Workshop" January 2025

448G Host Channel Topologies Analyzed

Building on initial feasibility studies, Synopsys expanded its analysis by working with 448G channel models from Samtec. These included 12 distinct channels spanning three topologies—CableBP, CurrentGenCPC, and NextGenCPC–across four different lengths.

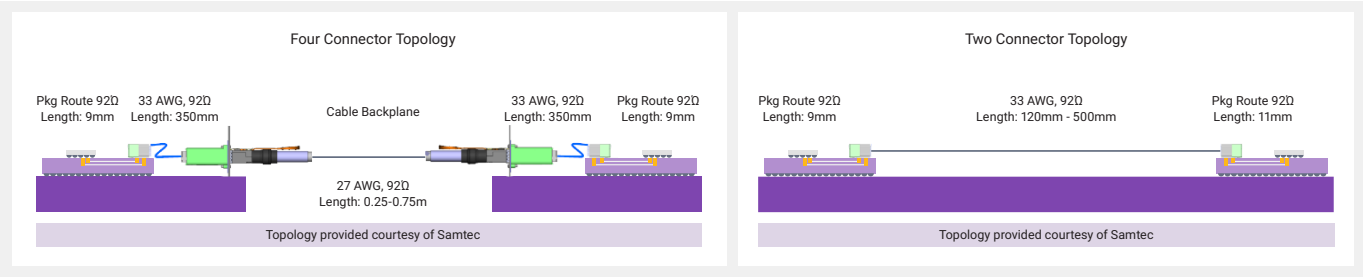


Figure 3: 12 (3 topologies with 4 lengths per configuration) Channels Provided by Samtec

Synopsys used these models to simulate 448G SerDes behavior, assessing system margins and confirming that these channels are sufficient to support early 448G deployments, especially in AI-driven compute clusters

The analysis focused on S-parameter frequency responses, particularly the insertion loss profiles and notches near the Nyquist rate (~112 GHz for PAM4, ~90 GHz for PAM6). These features are critical in determining whether the channels can maintain sufficient signal integrity and eye margin at ultra-high data rates.

These SerDes simulations across all 12 channel models using both PAM4 and PAM6 modulation schemes showed:

- **CableBP channels** exhibited moderate insertion loss and were generally well-suited for PAM4, benefiting from its lower Nyquist frequency and relaxed channel requirements
- **CurrentGenCPC channels** showed tighter constraints near Nyquist, posing greater challenges for PAM4 without aggressive equalization or design tuning
- **NextGenCPC channels** demonstrated improved responses over CurrentGenCPC, indicating promising potential for supporting PAM4 at 448G with optimized DSP architectures

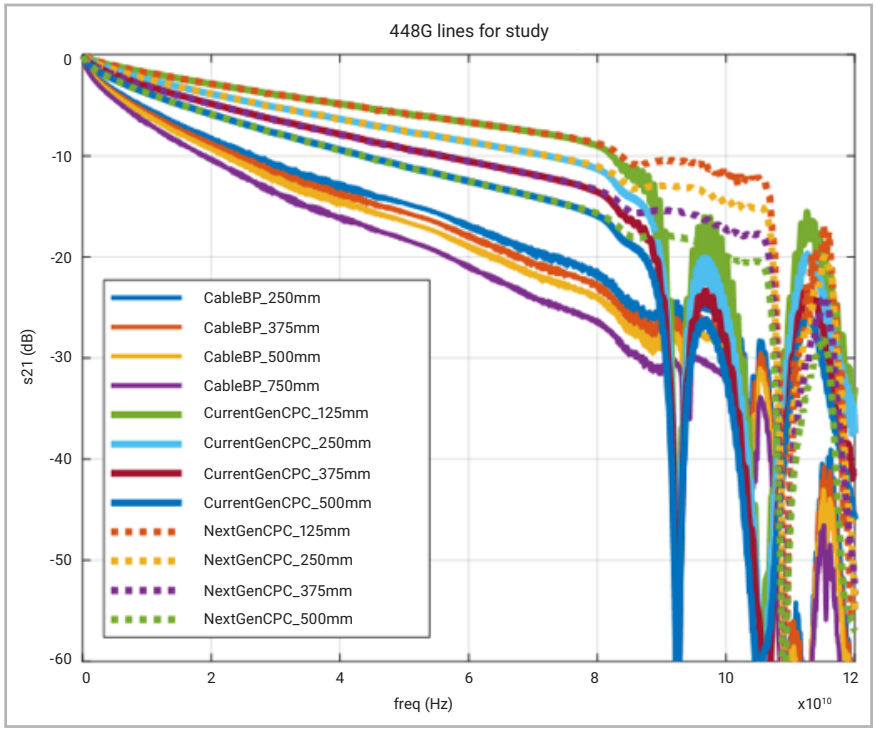


Figure 4: Synopsys channel result simulations

448G Analysis Summary

To evaluate the feasibility of 448G across real-world channel conditions, Synopsys performed comprehensive SerDes simulations using the 12 Samtec-provided channel models. These analyses assessed system margin under increasingly realistic impairments—including crosstalk (Xtalk), jitter, random noise, and Level separation mismatch ratio (RLM)—across both PAM4 and PAM6 modulation schemes.

The simulations were organized into three configurations:

1. **Xtalk Only:** Evaluates ideal baseline system margin assuming only adjacent-channel interference
2. **Xtalk + Jitter + Noise:** Adds dynamic impairments such as timing variation and receiver thermal/electrical noise
3. **Xtalk + Jitter + Noise + RLM:** Incorporates level separation mismatch ratio (RLM), which is critical for higher-order modulation receiver performance

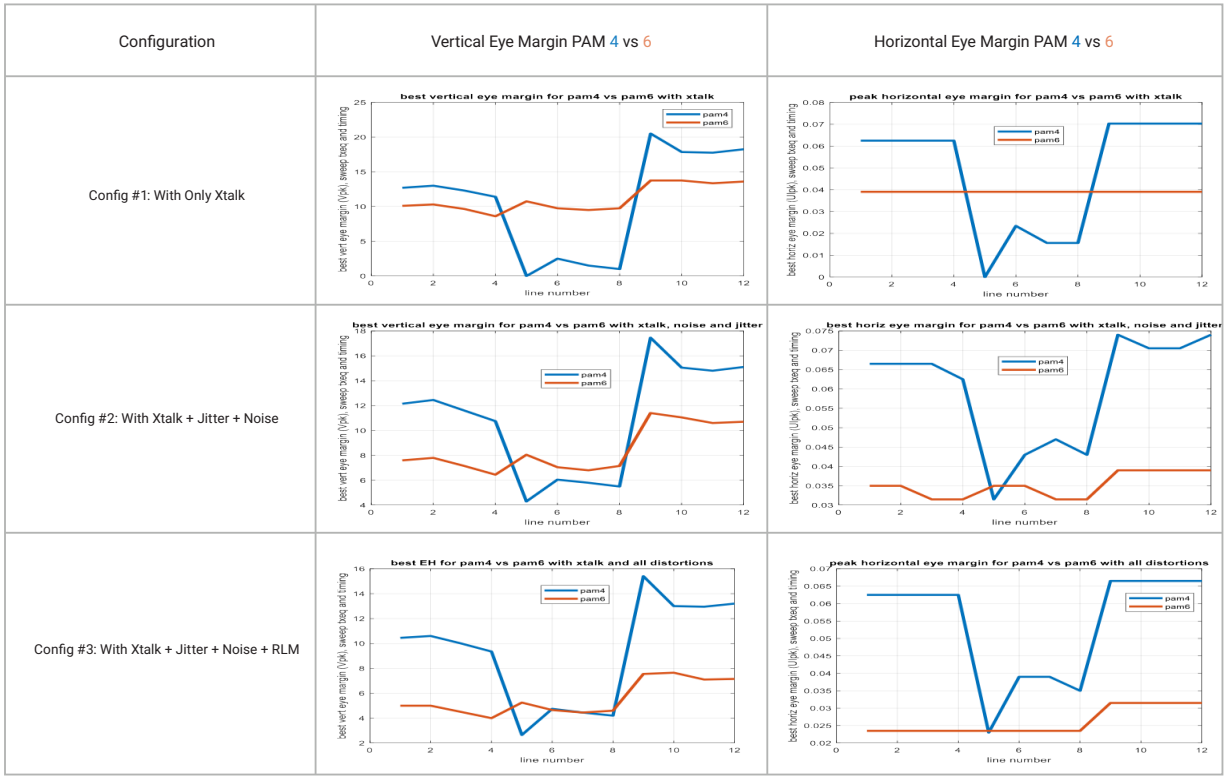


Figure 5: Evaluating Eye Margins Across Configurations

The results show that:

- Both PAM4 and PAM6 are technically viable across the studied channels, though their tolerances differ under various impairment combinations
- PAM6 benefits from lower Nyquist frequency and is more forgiving in terms of channel loss, especially under Xtalk-dominated conditions
- PAM4, while more power-efficient and compatible with existing optics, demands tighter transmitter control—particularly when RLM effects are introduced
- Designers retain meaningful tuning flexibility (“knobs”) in both transmitter and equalizer design to close margin gaps under realistic system constraints

DSP Architecture Enhancement for PAM6

Supporting PAM6 at 448G introduces a increase in DSP complexity compared to PAM4. While PAM6 offers advantages like lower Nyquist frequency and extended electrical reach, these benefits come at the cost of architectural overhead in the receiver and transmitter paths. Some of the added complexity stems directly from the nature of PAM6 modulation. With six amplitude levels instead of four, the DSP must handle finer voltage distinctions between levels. This leads to greater demand for silicon area and power within the receiver.

Impact on MLSD and DSP Scaling

The Maximum Likelihood Sequence Detector (MLSD), a core block in high-performance SerDes DSP, becomes notably more complex with PAM6. MLSD uses a trellis-based algorithm to evaluate symbol transitions, and the number of trellis states increases with modulation order:

- PAM4 requires support for 16 transitions
- PAM6 expands this to 36 transitions, demanding a more complex MLSD implementation

Table 2. PAM4 vs PAM6

Block	PAM4	PAM6
MLSD	16 transitions (4x4)	36 transitions (6x6)
Unrolled DFE	16 comparators (4x4)	36 comparators (6x6)
DLEV (Data levels)	4 Levels	6 Levels
Detector	2 bits	3 bits (for 6 levels)

Despite this, modern Viterbi-based MLSD architectures—already proven at 224G—are now being adapted for 448G. Innovations in latency reduction and hardware efficiency make MLSD feasible for PAM6, though it remains a challenging design task.

ADC and DAC Resolution Considerations

PAM6 also places stricter demands on data converter resolution:

- The receiver ADC may require an additional bit of resolution compared to PAM4 to resolve the finer step sizes between amplitude levels. However, if the system is dominated by channel noise rather than ADC quantization noise, this additional bit may not provide meaningful benefit and could be avoided to save power
- The transmitter DAC is more likely to require an extra bit to maintain coefficient resolution across the wider dynamic range of PAM6, especially for maintaining signal linearity and dynamic range

While PAM6 allows relaxed channel requirements and longer reach, it does increase DSP and converter complexity. Implementers must evaluate whether these trade-offs—higher power, area, and design effort—are justified for their specific use cases, particularly in power-constrained AI and HPC environments.

PAM6 is feasible at 448G, but it requires thoughtful system design and architectural innovation to remain within acceptable power and performance envelopes.

View and Directions—As of Q4 2025

The transition to 400 Gbps/lane Ethernet represents the next leap in SerDes and system design complexity, requiring coordinated innovation across host architecture, channel design, modulation strategy, and DSP implementation. This whitepaper explores how both PAM4 and PAM6 are viable modulation schemes for 400 Gbps/lane Ethernet deployments, each presenting distinct trade-offs in terms of signal integrity, reach, DSP complexity, and power consumption. While PAM4 benefits from compatibility with existing optical systems and power efficiency, PAM6 offers greater margin in lossy electrical channels at the cost of higher DSP and silicon complexity.

Through detailed simulations using real 448G channel models, Synopsys has shown that feasible operating margins can be achieved using either modulation approach with appropriate equalization and architectural tuning. The choice between retimed and unretimed host systems also plays a critical role in balancing performance and power across different deployment scenarios—especially in AI-driven scale-up and scale-out compute networks.

As the ecosystem progresses, designers will need to navigate a broad set of trade-offs tailored to their application, power envelope, and latency constraints. Synopsys continues to support this evolution by contributing simulation results, silicon data, and architectural guidance to standards bodies and industry working groups.

To explore how these physical-layer design choices align with broader system requirements—including network scaling trends, industry standardization efforts, and deployment priorities across AI and HPC workloads—read our companion whitepaper: Enabling the 448G Era: System Architecture and Standards for Next-Gen AI Networks.
It provides the strategic context for 448G, highlighting the industry forces driving the need for next-generation connectivity and how organizations are aligning around a common set of requirements. Together, these two whitepapers offer a complete view of the path to practical, performant, and standards-aligned 448G Ethernet solutions.

SYNOPSYS®
©2025 Synopsys, Inc. All rights reserved. Synopsys is a trademark of Synopsys, Inc. in the United States and other countries. A list of Synopsys trademarks is available at http://www.synopsys.com/copyright.html. All other names mentioned herein are trademarks or registered trademarks of their respective owners.
09/17/25.SNPS1740185146-Designing-for-448G-Modulation-DSP-WP.