

For PAM4
Data Center
Applications

Tri-Edge™

Enabling Your Analog Optical Switch Strategy



Executive Summary

In the coming years, data center traffic is predicted to continue its immense growth. To support the higher demand, data centers will require networking at higher speeds with higher capacity servers, switches and interconnects using PAM4 signaling.



For optical fiber interconnects, next generation analog modules have emerged to meet these advanced needs. When considering the move to PAM4, data center professionals need to understand whether Digital Signal Processing (DSP)-based optical modules offer the right value proposition based on cost, power, latency and performance. For systems using 200, 400 or 800G, this

paper examines the key technical advantages and value proposition for Semtech's Tri-Edge™ analog technology for optical modules, and discusses how analog solutions for PAM4 reduce power, latency and cost.

As massive growth continues in Cloud data centers from both private and public Clouds, workloads and

compute instances are growing at annual rates exceeding 20 percent. New Cloud data centers are being built to handle traffic from Cloud consumption, but according to the Cisco Global Cloud Index, 85 percent of data center traffic does not go out to users. 71 percent of this traffic is East-to-West within the data center and 13 percent is data center-to-data center. Only 14.9 percent is data center-to-user.

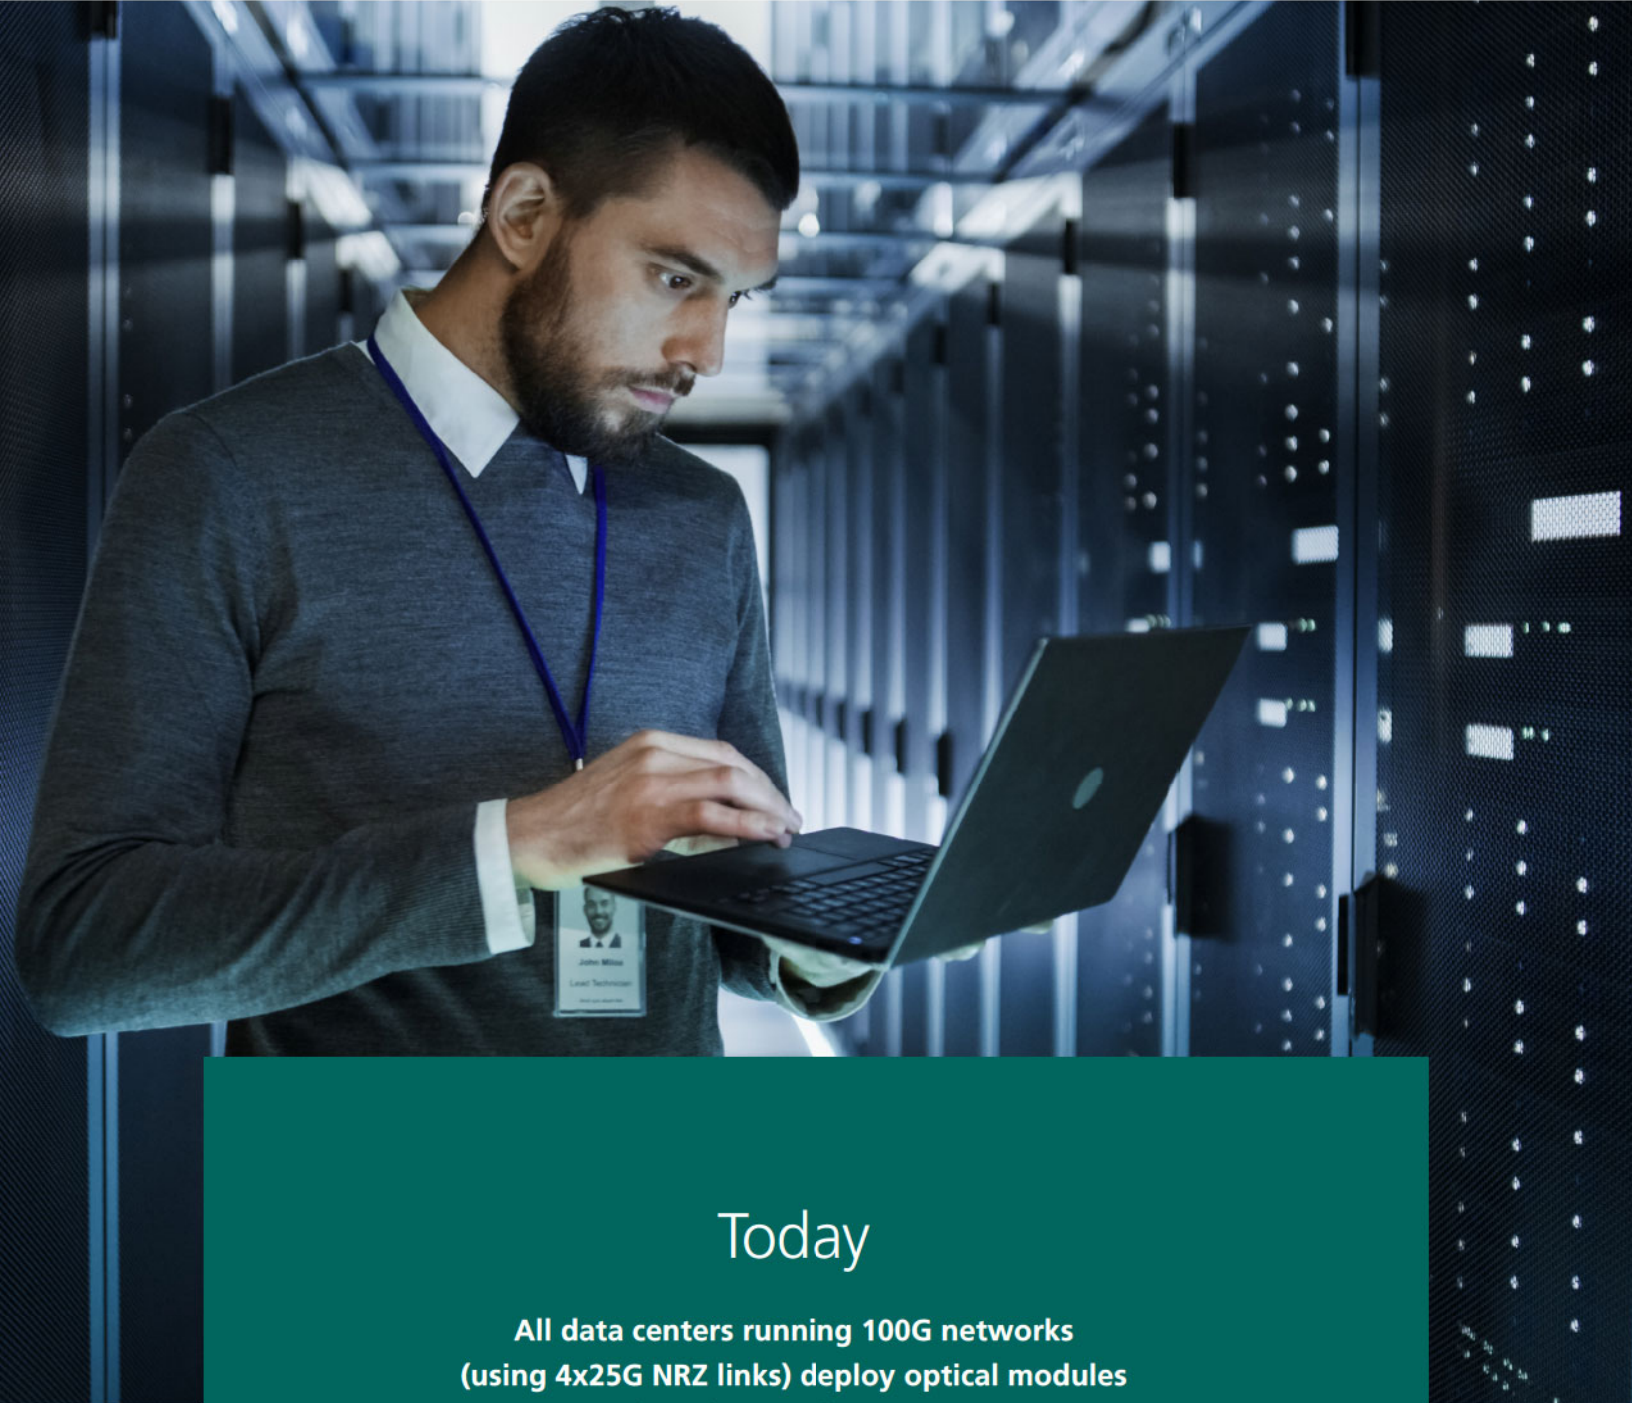
Massive within-data center traffic is multiplying the need for additional bandwidth and faster networking interconnection speeds. Most major data centers are in transition to 200 or 400G, but until recently the only choice was DSP-based optical module solutions that consume high power, suffering from higher latency and costs.

Semtech's Tri-Edge PAM4 clock and data recovery (CDR) technology is the first low power solution available on the market for optical communication links that addresses

multiple markets, from data centers and high performance computing (HPC), to 5G wireless fronthaul/midhaul/backhaul, with the vast majority of links being within 100m to 10km. This paper explores the benefits of the Tri-Edge CDR platform in the data center.

Semtech's Tri-Edge PAM4 platform builds on the widely successful ClearEdge® non-return-to-zero (NRZ) product platform. Semtech's Tri-Edge platform supports multiple optical modules for reaches from 100m to 10km. To achieve this, the Tri-Edge platform supports optical modules and active optical cables (AOC) using vertical-cavity surface-emitting lasers (VCSEL), directly modulated lasers (DML) and electro-absorption modulated lasers (EML). Tri-Edge enables optical module solutions meeting IEEE compliance and interoperability standards that are also fully compliant with the Open Eye MSA specifications for short reach (SR) and long reach (LR) links.





Today

All data centers running 100G networks
(using 4x25G NRZ links) deploy optical modules
based on analog CDRs.

Advantages to Analog Solutions

Analog Was the NRZ Solution Until PAM4

It took the optical industry about 10 years to transition from 10 to 100G. Today, all data centers running 100G networks (using 4x25G NRZ links) deploy optical modules based on analog CDRs. The highest volume deployments are for optical transceiver supporting links less than 100m (called SR or Short Reach, based on VCSEL technology). Semtech's ClearEdge CDRs have a major market share of CDRs in the SR optical module market due to the devices' superior power consumption, features that enable low module cost and a high level of integration.

The PAM4 market, which will take the industry to 200, 400 and 800G, has until recently used DSP-based module solutions; beginning with its release in 2018. Analog optical PAM4 CDR module solutions did not enter the market until 2019 with Semtech's Tri-Edge product line.

Driving the industry towards higher data rates is the increase in data processing inside the data center, requiring massive interconnections from thousands of simple network cards connecting servers through switches transmitting data across the networks. As silicon switch suppliers move at astonishing rates to release faster chips supporting higher data rates and capacity, data center customers are also weighing the factors of cost, power and performance of the optical interconnects.

ClearEdge Paved the Way for Faster Speeds

Servers, which process data, are connected to other servers and storage via Top of Rack (TOR) and Tier 1 switches. Increasing the link capacity and switch speed enables more data to be transferred between servers, and between servers and storage. As switch chips get faster, the bitrate of the links also must increase to match the needs of the servers. Semtech's ClearEdge CDR solutions are widely deployed today in 100G (4x25G NRZ) data center optical interconnects. As data center servers, switches and interconnects move past 100G, PAM4 signaling is used to double this bitrate versus NRZ. At the time of standardization of PAM4 optical link protocols, the low performance of the optics of the time required DSP-based solutions for these links.

Despite their fast interface speeds, DSPs use parallel data paths internally to process the converted analog signals and digitally process them before converting them back to analog at the output. The use of parallel processing adds power to the DSP solution as well as latency to the link. In contrast, analog ICs, including Tri-Edge, process the signal at the signal bitrate which lowers the chip's power and minimizes the latency. Semtech's analog technology also enables integration of the laser driver and transimpedance amplifier (TIA) which minimizes the interfacing power requirements and helps to maintain signal integrity. ClearEdge focused on NRZ bit rates and operates at speeds up to 25G. Tri-Edge is focused on higher speed PAM4 which operates at 50 and 100G per channel.

Different Fibers for Different Reaches

Each data center operator has its own architecture to optimize performance for the applications deployed. Generally, for the lowest tier of connections with a reach of below 3m it does not make sense to use fiber, optical products or active components. Copper cables make up most of these connections and have the lowest cost and lowest power.

Connections with a reach greater than 3m require optical cables and optical modules. The volume of optical interconnects are lower and costs are higher than copper interconnects. The second tier of connections generally use multimode fibers (MMFs), which are often used for short reach (SR) distances

(<100m at bitrates of 25G and higher).

Multimode fiber enables the use of the lowest cost optics, but the reach is limited by the modal bandwidth of the fiber. Semtech's ClearEdge ICs work in MMF NRZ applications and focus on the SR4 and SR2 channels to address up to 400m reaches at 10G NRZ, and up to 100m reaches for 25G NRZ. For longer reaches, single mode fiber (SMF) is used. This can either be deployed as single fibers or as parallel single mode (PSM) fibers.

Semtech's new Tri-Edge platform fits well into current data center fiber infrastructure for MMFs and SMFs, and enables implementations using 200 and 400G PAM4 for reaches from 100m to 10km.

What Is Tri-Edge?

Building on the success of the ClearEdge NRZ-based CDR platform, Tri-Edge is a new approach to PAM4. It is a platform of products addressing 200, 400 and future 800G applications for data centers, AOCs, and wireless markets. Tri-Edge is based on analog CDRs rather than DSP.

Tri-Edge SR products are for multimode fiber links up to 100m, while the Tri-Edge LR platform enables 200 and 400G SMF solutions up to 10km, enabling the product family to cover all the reach needs of intra-data center links.

Tri-Edge CDRs are multi-lane signal conditioners with best-in-class sensitivity, output jitter and jitter tolerance while offering the lowest power and cost solution for optical modules including:

- DSFP
- QFSP56
- QSFP-DD
- OSFP
- SFP56
- SFP-DD
- 50G AOC
- 100G AOC
- 200G AOC
- 400G AOC









Why Tri-Edge?

Networking up to 800G necessitates choosing an optical module strategy with the flexibility to meet future data center demands.

Key technical differentiators that make Tri-Edge a proven choice for multi-lane signal conditioners include:

- Monolithically integrated analog PAM4 CDR with laser drivers and TIAs
- Compact, high-performance solution with fewer chips
- Low power
- Low latency
- Low cost

Tri-Edge offers a more targeted approach to networking when compared to current DSP solutions, focusing on lower power with the right level of performance and value to benefit both AOCs and modules.

	SEMTECH	COMPETITION
Typical Power <ul style="list-style-type: none"> • <3.5W @ 200G • < 7W @ 400G vs. 10W @ 400G for DSP solution 		
Latency <ul style="list-style-type: none"> • 0.3ns vs. 10-40ns for analog solution and >100ns for DSP solution 		
Performance <ul style="list-style-type: none"> • Can enable different VCSEL vendors (ongoing) • FEC error free and 0 packet drops over temp 		
Integrated Chipset <ul style="list-style-type: none"> • Lowest BOM cost solution available 		

Solving the Power Problem

Power is a large factor in optical module performance as typical applications demand a significant portion of the data center's overall power budget. Power caps exist in data centers and increases in power used for interconnects can reduce throughput, bandwidth and processing in the center. Lower power connections between servers and switches offer major operating cost savings to the data center operator and to end users when upgrading applications.

The majority of hyperscale data centers are transitioning from 100G NRZ modules to higher data rates using PAM4. Current DSP implementations consume much higher power than Tri-Edge and its SR chipset. Customers report up to 30 percent power savings versus competing DSP solutions. The Tri-Edge SR chipset enables sub 3.5W (200G) and sub 7W (400G) solutions. DSP 400G module solutions consume around 10W. This means Tri-Edge not only meets customer module power requirements but Tri-Edge also enables QSFP-DD and OSFP module form factor classes (3.5W for 200G and 7W for 400G) that DSP cannot.

Low power is also vital in Artificial Intelligence (AI)-centric data centers as processing consumes huge amounts of power, leaving less for data interconnects (including optical modules). Machine Learning (ML) needs high throughput, with low power and low latency for processing. Tri-Edge's capabilities make it an ideal solution for these applications.

Integration is a major factor in Tri-Edge's low power capabilities. Having an integrated IC with both CDR and PMD (VCSEL, DML or EML drivers, or TIA) in a monolithic solution enables the lowest power solution in the market as it eliminates power associated with chip-to-chip interfaces in the optical module. Integration means fewer chips, fewer interfaces, lower cost, and a smaller footprint.

This integrated analog approach means the module only requires two chips, one integrated IC with a laser driver and one IC integrated with a TIA. Integration enables a much simpler, more elegant solution compared with power hungry DSP modules. Typical DSP module architectures need high power to drive separate interfaces between discrete chips. The high power can create thermal gradients/issues requiring more thermal management.

Semtech's ClearEdge product line pioneered the integration of the laser driver and TIA into a single chip to enable low power and other benefits. ClearEdge continues to lead the market today. Tri-Edge brings these same benefits to existing PAM4 solutions and to future applications.

In a recent webinar on Lightwave Online called "System Evaluation of On-Board Optics," the COBO members presented data showing Tri-Edge based SR8 modules operating at 35 percent less power than comparable DSP based modules.

Integration Creates Lowest Latency Solution Available

In supercomputing data centers with compute and data intensive applications (e.g. HPC, cloud, and AI), data processing time is critical. As the complexity of mathematical modeling increases, the need for additional compute power grows.

In environments that strive to maximize compute throughput and reduce server idle time, a low latency network is critical, and existing interconnect delays can reduce a supercomputer's overall performance significantly. In these parallel computing environments, interconnect latency has a large effect on overall performance and the lowest latency is required.


Cable Latency Multiplication

3-level fat tree (extract shown)


- 6 cable hops

Latency per unit


- Nic 150ns
- Switch 130ns
- Flight Time 6ns/m
- Optical Receiver 20ns vs. 120ns



System Latency – Open Eye optics



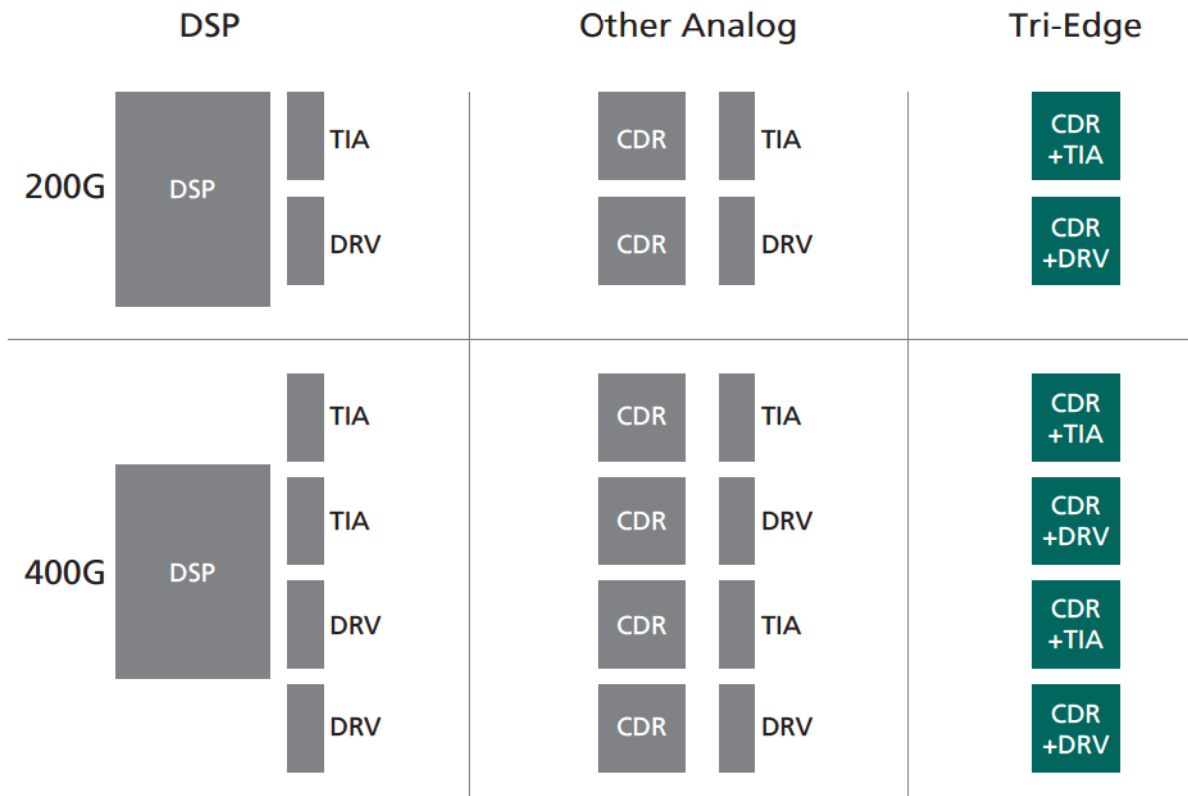
System Latency – DSP based



Tri-Edge’s integrated CDR chipsets offer unmatched low latency – among the lowest in the industry. This is considerably lower than the typical DSP solution latency of greater than 100ns.

The key to low latency is Tri-Edge’s analog signal integrity solution where the PAM4 signal is equalized and retimed in real time as the signal passes through the Tri-Edge chip. With Tri-Edge there is no need for an analog to digital convertor (ADC) to digitize the incoming signal followed by digital signal processing algorithms and then digital to analog conversion (DAC) to resend the signal, as is done in DSP-based solutions.

The Tri-Edge platform provides this low latency while also monolithically integrating the optical transmit driver (LDD) and the optical receive TIA for best performance and lowest power. DSP competitors need at least three ICs for 200G (Tx/Rx DSP, Tx LDD, and Rx TIA) and five for 400G (Tx/Rx DSP, 2x Tx LDD, and 2x Rx TIA) based on 50G per lane signaling. Other competitors’ analog PAM4 discrete solutions require four ICs for 200G (Tx CDR, Tx LDD, Rx CDR, and Rx TIA) and eight ICs for 400G (2x Tx CDR, 2x Tx LDD, 2x Rx CDR, and 2x Rx TIA) based on 50G per lane signaling. Along with this increased complexity comes higher cost and higher latency versus Tri-Edge.




Lower Total Solution Cost

Tri-Edge offers a low cost solution by enabling use of lower cost optics and offering fully integrated transmit and receive chips that occupy a small footprint. Tri-Edge also enables re-use of Semtech's established ecosystem of cost efficient manufacturing. The Tri-Edge platform does not require use of expensive advanced node complementary metal-oxide semiconductors (CMOS) because there is no need for high speed ADC, DAC and signal processing, as done in DSP-based solutions. The product development cost of advanced node CMOS DSPs is also much higher than Tri-Edge due to expensive mask set costs for the advanced node CMOS. These increased development costs are typically amortized into the cost of the DSP.

In addition to the lower development and material costs with Tri-Edge, the proprietary laser compensation technology enables the use of low cost lasers for PAM4 applications. For example, enabling 25G NRZ-class electro-absorption modulators (EMLs) provides significant cost savings over high-bandwidth 25G PAM4-class EMLs. This allows 200 and 400G solutions to leverage lower cost 100G (4x25G) optics that are in use in high volumes today, enabling shorter design cycles and easier sourcing of components.

Tri-Edge,
proprietary laser
compensation technology
enables the use of low
cost lasers for PAM4
applications.



Open Eye MSA: Easing the Transition

Semtech is leading the way to enable analog PAM4-based solutions in the data center market and creating a complete analog PAM4 ecosystem with the Open Eye MSA. Semtech, along with other leading companies who recognize the value of analog capability, formed the Open Eye consortium in 2019. This is a multi-source agreement (MSA) to accelerate the adoption of PAM4 and enable lower cost modules from 50 to 400G or more in data centers and other applications.

Previously, standards committees developed link and switching standards based on the available DSP solutions, which required significant equalization in the links. As optical components have improved, the industry recognized that PAM4 using lower equalization analog solutions as much simpler and providing the right value point for what data centers need: lower cost modules with simplified, robust and repeatable compliance requirements.

The Open Eye MSA leveraged the advancement of optical performance over the past few years to provide a technology-independent specification enabled by DSP and/or analog PAM4 solutions. Low cost optics, components and systems providers needed to define an interoperability specification to

enable multiple supply sources for all components. The Open Eye MSA has developed industry standard optical interconnects allowing for a wide variety of architectures and technologies and standardized specifications to make compliance testing less complex and improve interoperability. In addition, the Open Eye MSA specification utilizes well-proven eye mask metrics for transmit and receive that enable interoperability. This eye mask specification is easily implemented in test equipment and has been an established standard link specification metric in NRZ systems. The Open Eye MSA also creates an ecosystem around the implementation of PAM4 where end users can choose from multiple sources of technology, including low power analog suppliers.

Semtech is the first analog PAM4 provider to be in production. Its Tri-Edge products passed Forward Error Correction (FEC) and customer qualification testing, and are leading the industry with implementations. Together with the Open Eye MSA, which now has more than 35 members, there is strong interest in the continued success and growth of analog-based PAM4 optical solutions.



Future Ready

Tri-Edge will also enable data centers to move to higher data rates (from 100G or above to 800G), and the future AI, supercomputing and 5G wireless infrastructure markets will also benefit from its key advantages. Tri-Edge CDRs consist of a family of solutions that address a data center's optical interconnect needs from 100m up to 10km while offering best in class power, latency and cost. SR solutions are in production now, while LR solutions for single mode fiber are sampling and are expected to be in production by the end of 2020.

The flagship Tri-Edge CDR products currently in production are the GN2558 and GN2559 to enable 200 (4x50G) and 400G (8x50G) AOCs/SR4/8 for the 70m OM3 and 100m OM4 reach markets. These products integrate the PMD functions and pave the way towards higher data rates.

The next two products in the Tri-Edge family will address the 50 and 100G SR market up to 100m, enabling low cost VCSELs. For single mode fiber applications, the Tri-Edge platform will provide fully integrated analog PAM4 CDRs with EML or DML drivers to enable the 200 and 400G 2km reach and TOR-to-server switch market. Collectively, the Tri-Edge portfolio addresses the entire range of switching interconnect needs within a data center today and will enable PAM4 solutions as data centers move to 800G and beyond. Tri-Edge solutions are also expected to find implementation in the fast growing wireless infrastructure markets, including 5G markets. Typically driven by mobile carriers, Tri-Edge solutions will play a key role in connecting data from mobile phones, autonomous vehicles and other wireless devices from towers to data centers to meet data processing and bandwidth needs.

Semtech Data Center Multi-Mode Offerings

GN2558 Quad PAM4 CDR + Linear VCSEL Driver

GN2559 Quad PAM4 CDR + Linear TIA



- Enables 200G (4x50G) and 400G (8x50G) AOCs/SR4/8
- Reach: 70m OM3, 100 OM4



GN2538 Dual PAM4 CDR + Linear VCSEL Driver

GN2539 Dual PAM4 CDR + Linear TIA

- Enables 50G (1x50G) and 100G (2x50G) AOCs/SR2
- Reach: 70m OM3, 100 OM4



Conclusion

Semtech's Tri-Edge technology offers the only analog CDR solution for optical modules capable of meeting the low power, low cost requirements needed for data center PAM4 optical interconnects. Tri-Edge also offers significant latency improvements over DSP which are key to HPC and AI data centers demanding the lowest latency. The products are fully compliant to the Open Eye MSA and address a full range of needs for a data center, from 500m to 10km, SR, TOR and Tier 1, in a range of laser and module types.

Products enable interoperability and fast module bring up, and are currently in production for MMF SR applications SR4, SR8, AOC applications with the GN2558/59 chipset. This will be followed by the GN2538/9 set for MMF SR applications, SR2 and SFP56 and Tri-Edge chipsets for single mode, long reach (LR) modules. Multiple module customers are designing products around Tri-Edge and completing full data center testing.

About Semtech

Semtech Corporation is a leading supplier of high performance analog and mixed-signal semiconductors and advanced algorithms for high-end consumer, enterprise computing, communications, and industrial equipment. Products are designed to benefit the engineering community as well as the global community. The Company is dedicated to reducing the impact it – and its products – have on the environment. Internal green programs seek to reduce waste through material and manufacturing control, use of green technology and

designing for resource reduction. Publicly traded since 1967, Semtech is listed on the Nasdaq Global Select Market under the symbol SMTC. For more information on Semtech's optical products, including the Tri-Edge portfolio, visit: [semtech.com/products/signal-integrity](https://www.semtech.com/products/signal-integrity)

Transform Your Optical Applications



Tri-Edge™
Ultra-low Latency
Signal Integrity Solutions

SEMTECH

Signal Integrity Products

Learn more about its leading optical networking products for data centers, 5G networks and more on the Semtech website, and visit the videos page to watch the Tri-Edge video.

semtech.com



Semtech Corporation
200 Flynn Road
Camarillo, California 93012
SEMTECH.COM



SEMTECH.COM/PRODUCTS/SIGNAL-INTEGRITY