BCS Allegient

From Defense Research to Market Disruption: The Trillion-Dollar Potential of Advanced Laser Diode Technology

TELECOMMUNICATIONS AND DATA CENTERS

October 2025

From Defense Research to Market Disruption:
The Trillion-Dollar Potential of Advanced Laser Diode Technology

TELECOMMUNICATIONS AND DATA CENTERS

By: Charles D'Onofrio

Executive Summary

Artificial intelligence (AI) is reshaping data center infrastructure, driving unprecedented demand for ultra-high-performance optical connectivity. As AI model training grows in scale, data center computing resources must be interconnected with extremely low latency, creating strong demand for advanced optical transceivers. The U.S. optical transceiver market, valued at $13.6 billion in 2024 and growing at 13% annually[1], is becoming a critical enabler of AI and edge computing.

Yet today's short-reach transceivers face four fundamental limitations: diode efficiency capped at ~40%, reach constrained to 100–500 meters, high power consumption, and thermal degradation that forces five-year replacement cycles. These barriers increase operating costs and restrict distributed computing architectures.

Next-generation diode technologies – including photonic crystal surface-emitting lasers (PCSELs) and orbital angular momentum (OAM) multiplexing – could eliminate these bottlenecks. Higher efficiency lasers reduce heat, extend transceiver lifespans, and enable multi-kilometer reach, while OAM multiplexing and faster modulation unlock higher data rates. Together, these advances could transform hyperscale AI infrastructure and enable edge deployments not possible today.

A bottom-up analysis of the market potential offered by these technologies shows that between 2026 and 2040:

- Performance gains (speed + range) could unlock between $1 trillion and $8 trillion in U.S. market value.
- Reliability improvements could unlock another $1 trillion, even without performance gains.
- Electricity savings due to improved efficiency are negligible compared to acquisition costs.

Breakthrough diode technology represents a massive opportunity, with speed, range, and reliability emerging as the dominant value drivers.

Background: Laser Diodes in Data Centers and Telecom

Data centers house thousands of computers and memory storage systems, providing the computational power required to train large AI models. These systems exchange massive amounts of data at high speed, and as AI scales, connecting network computing resources has become the primary bottleneck for training performance. Optical transceivers are the hardware components that facilitate this connection.

Optical transceivers convert electrical signals into pulses of light for transmission across fiber optic cables. Each transceiver contains an optical engine, which performs the core functions of conditioning, transmitting, receiving, and processing optical signals. The optical engine consists of a transmitter, which converts electrical signals into optical pulses, and a receiver, which converts incoming optical pulses back into electrical signals. Transceivers are typically pluggable modules that insert into network switches or server interface cards (Figure 1), allowing operators to upgrade or replace individual units without changing entire switches or servers. Short-reach transceivers provide high-density connectivity within data centers, while long-reach transceivers connect data centers to telecommunication networks. State-of-the-art short-reach transceivers commonly operate at 800 Gbps (800G) and are limited to around 100 meters, while more expensive long-reach transceivers operate at 800G over kilometers.

Figure 1. *The left shows a pluggable optical transceiver (from FS.com). The finned end of the transceiver gets plugged into a switch and the black caps on the opposite end are removed to plug in optical fibers. The right shows a switch full of pluggable optical transceivers (from Unitekfiber.com).*

Training and hosting large AI models requires enormous computational power. Thousands of computing resources communicate with each other and memory storage through optical transceivers, and the speed of these connections, not just processor performance, limits training acceleration[1] (Figure 2). Within the transceiver, the transmitter converts electrical signals into optical pulses using a laser diode. Current laser diodes are limited by modulation speed, which determines how fast they can pulse, and by efficiency, converting only about 40% electrical input into optical output. The remaining 60% is lost as heat, increasing cooling requirements and constraining both performance and lifespan. Improved laser diodes with faster modulation and higher efficiency would increase data throughput, reduce power and cooling needs, and enable the extended-reach connectivity required for edge computing applications[2].

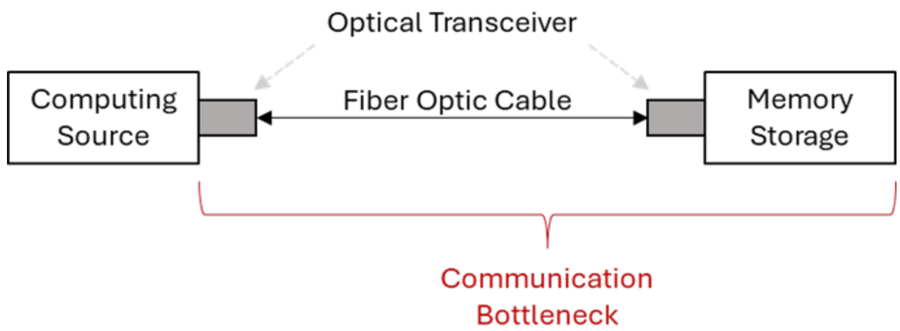

Figure 2. *It is the speed at which data is communicated that inhibits training AI.*

With this understanding of transceiver function and the role of laser diodes, we can now examine the engineering details of a representative short-reach system. The 800G optical transceiver illustrates how multiple laser diodes are arranged and modulated to achieve high data rates while operating within the thermal and efficiency limits discussed above. Studying this system provides a concrete baseline for evaluating the potential gains from next-generation laser diode improvements.

The 800G Breakdown: Engineering Analysis of Today's Optical Workhorses

800G Architecture and Performance

An 800G short-reach transceiver costs around $800 and contains eight laser diodes that are each optically coupled with eight unique fibers, typically multimodal fibers (MMF). The laser diodes generate light in the infrared regime at 850nm. Each laser diode can pulse utilizing 4-level Pulse Amplitude Modulation (PAM4), where pulse amplitude can represent a symbol of either 00, 01, 10, or 11. A single Vertical Cavity Surface Emitting Laser (VCSEL), which is the type of laser diode typical of 800G transceivers, can send 2 bits of information per symbol at 50 GBd[1] for a data rate of 100G. With 100G per laser diode across eight unique laser diodes, the transceiver achieves an 800G signal.

Power Consumption and Efficiency Limits

An 800G optical transceiver typically consumes 16W of power, where 3.5W is devoted to the transmitter and 2.5W is devoted to the receiver. The remaining 10W is devoted to other supporting electronics, mainly Digital Signal Processing (DSP). Relevant componentry including power and heat dissipation is shown in Figure 3.[2] State-of-the-art VCSELs within 800G transmitters convert electricity into optical energy with a 40% power conversion efficiency[3] (PCE)[5]. The remaining 60% comes off as heat, which requires cooling for the laser diode to function as designed.

1 A gigabaud (GBd) refers to one billion symbol changes per second. It's a unit of symbol rate, which is the rate at which the signal representing data changes states. Data Rate = Baud Rate * Bits / Symbol

2 All diagrams and power and thermal estimates are derived from information in [2], [5], [6] and [7].

3 VCSEL PCE typical of high-data rate transceivers are not explicitly published, 40% comes from the highest performing commercially available VCSEL designed for LiDAR applications.

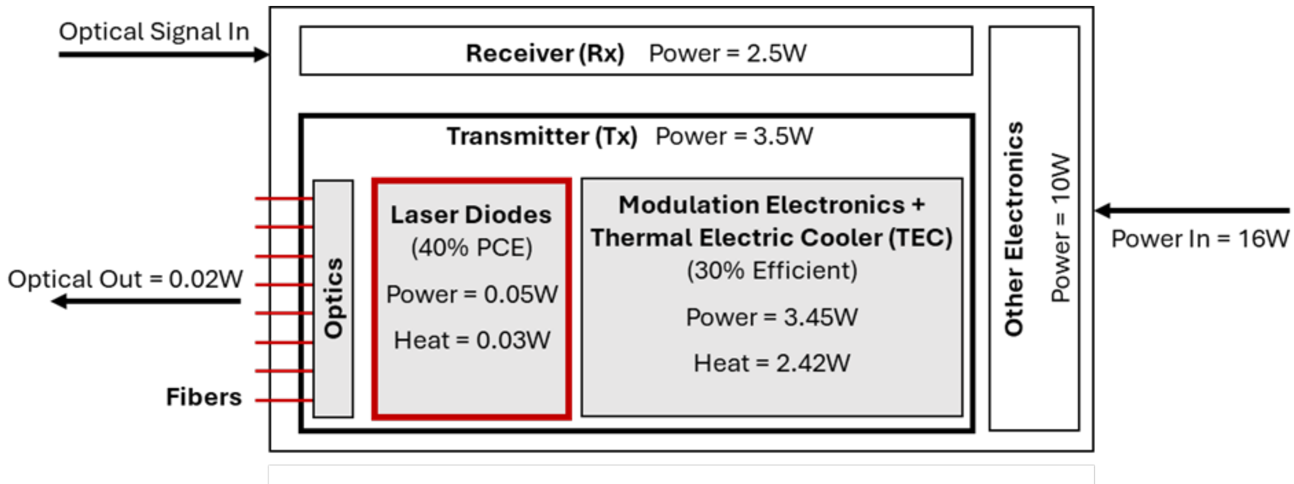

Figure 3. *A systems diagram of a state-of-the-art 800G optical transceiver showing power and heat dissipation within each subcomponent.*

VCSELs lose efficiency as they experience thermal cycles from operation. Over time, thermal degradation causes the laser to produce an optical signal too weak for the receiver to interpret, ending the transceivers' life. The typical service life of an 850nm 25GHz VCSEL in a data center optical transceiver is 88,000 operational hours,[9] which equates to five years at a 50% utilization rate.

Distance Limitations and Modal Dispersion

A key limitation of short-reach optical transceivers using VCSEL diodes is their restricted distance, typically less than 100 meters. This distance limitation stems from the inherent properties of VCSELs and their compatibility with multimode fiber. VCSELs emit light over a wide angle, making them difficult to couple efficiently into single-mode fiber. In multimode fiber, light travels along multiple paths of varying lengths, leading to modal dispersion. Modal dispersion, the difference in arrival times of these light paths as shown in Figure 4, distorts the signal which limits the achievable transmission distance.

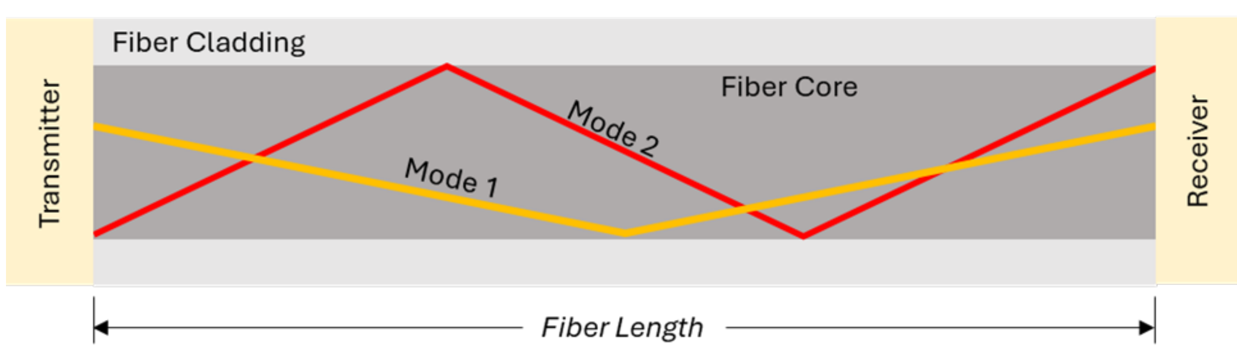

Figure 4. *Modal Dispersion occurs when mode 2, which travels a longer path distance than mode 1, arrives at the receiver too late behind mode 1 to be useful in creating a signal. This limitation occurs in the highest quality fibers at 150 meters.*

1.6T Evolution: Doubling Down on Brute Force

A handful of 1.6T (1600G) optical transceivers began commercial deployment in January 2025 priced around $3000 for single units. These 1.6T transceivers consume approximately 26W of power and use eight PAM4 modulated 100GBd Distributed Feedback (DFB) lasers to transmit 1310nm signals with an optical power of 100mW uncooled and 300mW cooled, achieving a reach of 500m[13] through single mode fiber (SMF).

These optical transceivers are a significant breakthrough for the telecom industry because they replace VCSEL diodes and multimode fibers for high-date rate communications with DFB laser diodes and single mode fiber. Nevertheless, power consumption, efficiency, reliability, and cost remain major concerns: the data rate doubled from 800G to 1.6T, but the energy requirement increased by 1.5x from 16W to 26W, the diode's efficiency remains at 40%, reliability likely suffers from the same thermal degradation induced failures mode within 5 years, and the cost has increased by 3.5x from ~$800 to ~$2800. This represents a brute force approach to faster communications.

The brute force approach is not indefinitely sustainable. Data center construction in developed markets, such as Data Center Alley in Ashburn, Virginia, is already limited by the amount of power they consume from the grid. Continually increasing optical transceiver power with each generation does not help overcome this limitation. What's needed is the opportunity that advanced laser diode technologies present: a breakthrough that increases speed and range while reducing power consumption and extending operational lifespans.

Technology Breakthrough: Beyond Brute Force Scaling

Two key advances in laser diode technology will overcome the brute force approach to increasing transceiver performance: (1) Orbital Angular Momentum (OAM) to increase data density, and (2) faster diode pulsing to increase baud rate coupled with improved efficiency to reduce heat generation. These approaches are under active investigation in academia and by a handful of startups, though their application within optical transceivers as a means of unlocking both performance and efficiency gains remain relatively unexplored.

Orbital Angular Momentum: Multiplying Data Capacity

OAM enables the creation of various far-field mode shapes at the receiver. The "far-field mode shape" is the pattern the light forms across the cross-section that hits the receiver, also called the transverse mode. Current transceivers use only one mode shape, but OAM technology can generate multiple distinct patterns as shown in Figure 5.

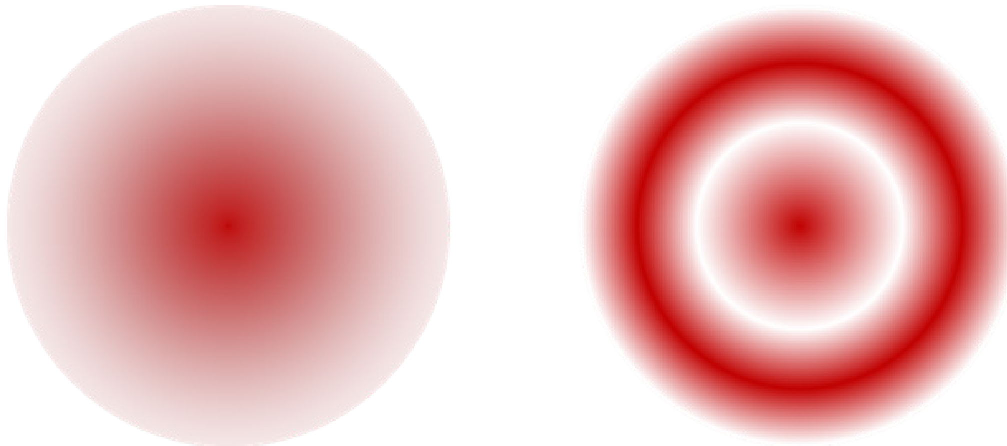

Figure 5. *Two different transverse mode shapes. The one on the left has a single peak in the center while the one on the right has two peaks -- one in the center and a ring that goes around the center.*

A laser diode that could selectively induce OAM could transmit more bits per second. If OAM enables two separate far-field mode shapes, this adds an additional bit to the existing PAM4 modulation scheme, generating eight possible symbols (000, 001, 010, 011, 100, 101, 110, 111), which increases the data transmission rate by 1.5x. Adding a third OAM mode would enable four bits per symbol, doubling the data transmission rate. This communication scheme is not used today because technology does not yet exist to selectively induce OAM modes at the laser diode.

Scenario A: Efficiency for Compatibility

Consider replacing the 40% efficient laser diodes in Figure 3 with 90% efficient laser diodes while reducing input power to maintain the same 0.02W optical output, as shown in Figure 6. This approach, which may be achievable by substituting VCSELs with Photonic Crystal Surface Emitting Lasers (PCSELs), would provide data center operators with a more efficient optical transceiver as a replacement compatible with current infrastructure. It would not require new optical fibers or new receivers.

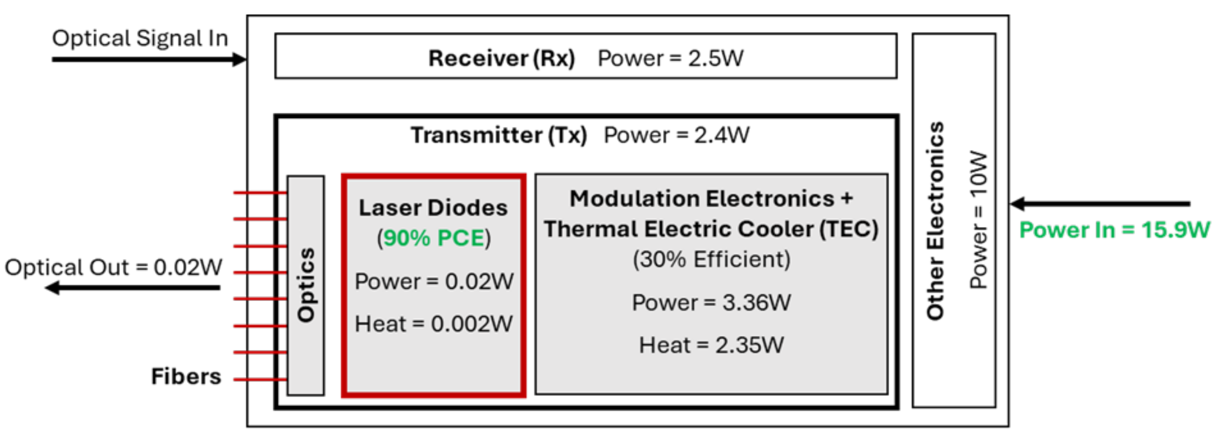

Figure 6. *The laser diode is 90% efficient and the optical output is conserved at 0.02W. Compared to the 800G baseline, this drops the power input required of the transceiver from 16W to 15.9W and reduces the heat produced by the transmitter by 0.09W.*

In this scenario, heat generated by each laser diode drops from 0.03W to 0.002W. This heat reduction decreases the temperature within the diode substrate by approximately 10°C. Per established VCSEL reliability models,[10] a 10°C temperature drop increases the service life of the diode by approximately 10x, practically eliminating thermal cycling as the primary failure mode in optical transceivers.

Scenario B: Efficiency for Performance

Alternatively, consider replacing the 40% efficient laser diodes with 90% efficient diodes while maintaining 16W input power and allowing optical output to increase to the greatest extent possible, which is 0.10W, as shown in Figure 7. This approach may be achievable by substituting DFB lasers with PCSELs, creating a technological leap beyond even 1.6T transceivers.

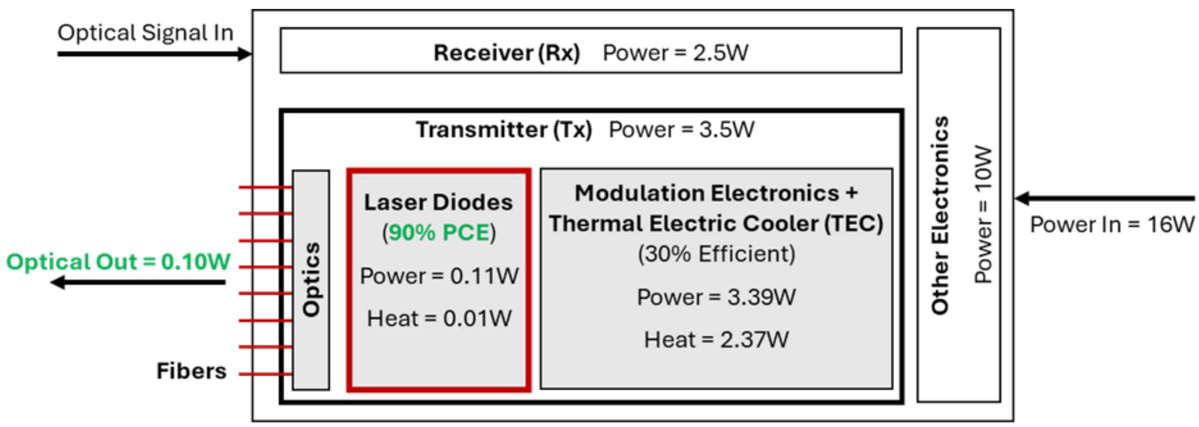

Figure 7. *The laser diode is 90% efficient and the input power is conserved at 16W. Compared to the 800G baseline, this increases the optical output to 0.10W and reduces the heat produced by the transmitter by 0.06W.*

The 5x increase in optical output power (from 0.02W to 0.10W) enables a 5x increase in range through single-mode fiber, while the improved efficiency still reduces heat generation compared to current technology. This breaks the brute force approach by delivering higher performance with better efficiency and reliability.

Combined Impact: Faster Cycling with Higher Efficiency

The concept of faster cycling is straightforward—if a diode can pulse at double its current speed, the baud rate doubles, however, the true breakthrough lies in cycling faster while simultaneously achieving higher efficiency. If realized, these innovations would not only address the physical limits of today's transceivers but also redefine the performance-to-cost tradeoffs that shape data center adoption.

Market Implications: How Much Are Increased Speed, Distance, Efficiency, and Reliability Worth

The previous section discussed two unique technological advances to overcome the brute-force approach to transceiver development: (1) utilizing OAM, and (2) faster, coupled with more efficient, diode modulation. This section evaluates the value proposition of these improvements, which includes sending data faster and farther, improving transceiver reliability, and reducing electricity costs per transceiver.

Methodology

The value proposition was assessed using a bottom-up methodology. A market survey conducted in July 2025 provided baseline performance-to-price data for current transceiver technologies. Price models were projected forward by assuming each transceiver generation follows similar pricing trends.

Demand curves were developed for each technology generation over time, allowing prediction of the product mix across years. These demand curves were then weighed by the forecasted price models to calculate an average transceiver sales price per year.

The installed base of transceivers in today's market was estimated using publicly available data and projected into the future based on market growth assumptions. Expected failures and replacements were modeled using Weibull reliability methods. Combining new demand with replacement demand produced annual unit sales forecasts, which were multiplied by the average transceiver price to estimate total annual revenue.

Market Survey

Quantifying the value of potential transceiver improvements requires a baseline understanding of how today's market prices scale with data rate and reach. To develop this baseline, a structured market survey of current optical transceiver pricing was conducted. The survey focused on four data-rate categories (100G, 400G, 800G, and 1.6T) and distinguished between short-reach and long-reach models.

- **Sample selection**: For each data-rate category, publicly available single-unit retail prices were collected from online distributors. Within each category, prices for three short-reach and three long-reach models were recorded when available.
- **Sample size**: This process produced six data points per category, except for 1.6T, where only two short-reach models were commercially available and no long-reach units were found. In total, the dataset contains 26 transceivers.
- **Scope and limitations**: Retail prices were chosen to provide a transparent, reproducible baseline, although actual bulk purchasing costs for hyperscalers are typically lower. Accordingly, the retail-based survey is used to establish relative price-performance relationships, while overall market growth and adoption are modeled separately based on historic demand patterns.

Price-Performance Relationships

Figure 8 shows the market survey results. No long-reach transceivers capable of achieving 1.6T data rates are currently available, so this category is omitted.

The relationship demonstrates exponential price scaling with data rate, suggesting buyers are willing to pay substantial premiums for speed: if the data rate doubles, the price quadruples. As expected, long-reach transceivers cost significantly more than short-reach transceivers for equivalent data rates and exhibit similar speed premiums.

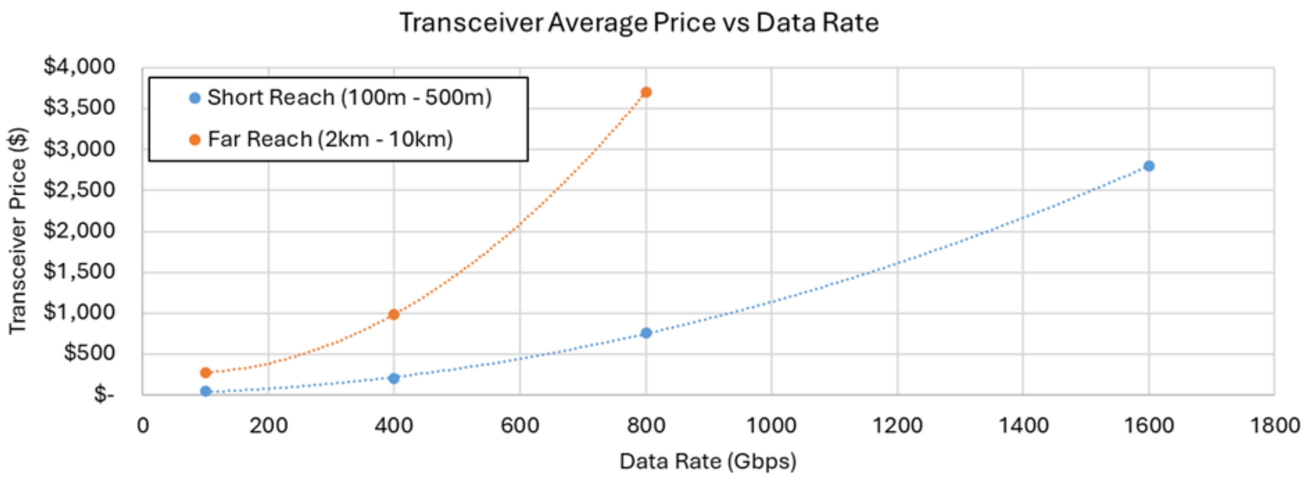

Figure 8. *Transceiver average sales price versus data rate for short-reach and long-reach transceivers.*

The market survey was extended by calculating the transceiver price per billion gigabits of data transmitted over its life, assuming 5-year operation at 50% utilization. Figure 9 shows the resulting trends. The linear relationships suggest buyers are not willing to pay a premium for life increase: if the amount of data sent over the transceiver life doubles, the price doubles.

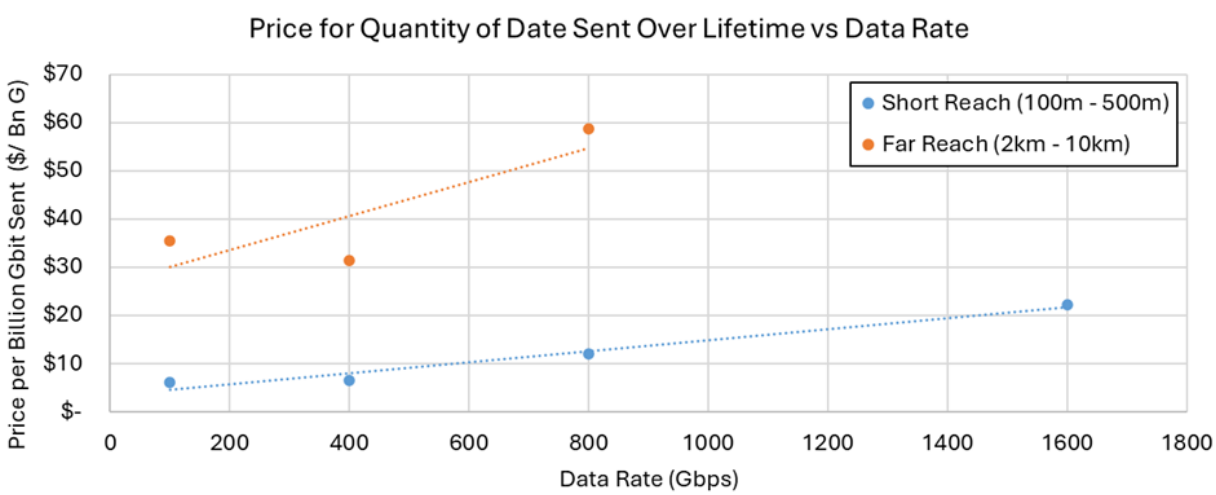

Figure 9. *The price per billion gigabits sent by a transceiver over the course of its life.*

Finally, a distance premium was derived to quantify buyers' willingness to pay for extended reach beyond short-reach capabilities. For example, extending the 100m range of an 800G transceiver costs $1.55 per additional meter. Figure 10 shows this result.

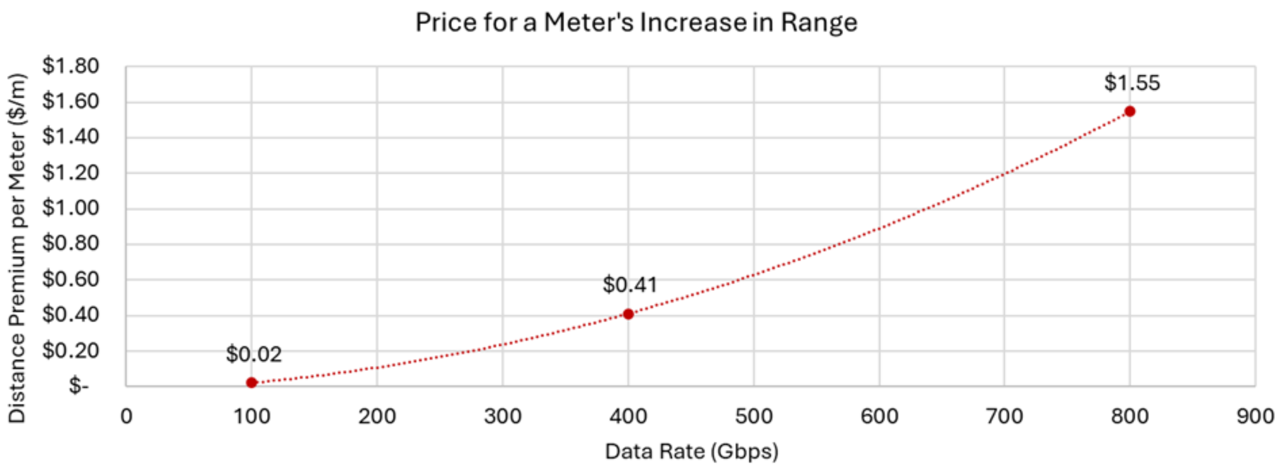

Figure 10. *The transceiver price increase for each additional meter of range at a given data rate.*

The market study shows that data rate and range are non-linear influences on transceiver prices. This information is useful for projecting a price range for transceivers with improved laser diodes.

Advanced diodes could enable a short-reach 3.2T transceiver matching the 500m range of current 1.6T systems. Building from the 1.6T price in the market survey, a 3.2T short-reach transceiver could sell in today's market for between $5,600 and $11,000 – between 2x and 4x the price of current 1.6T transceivers.

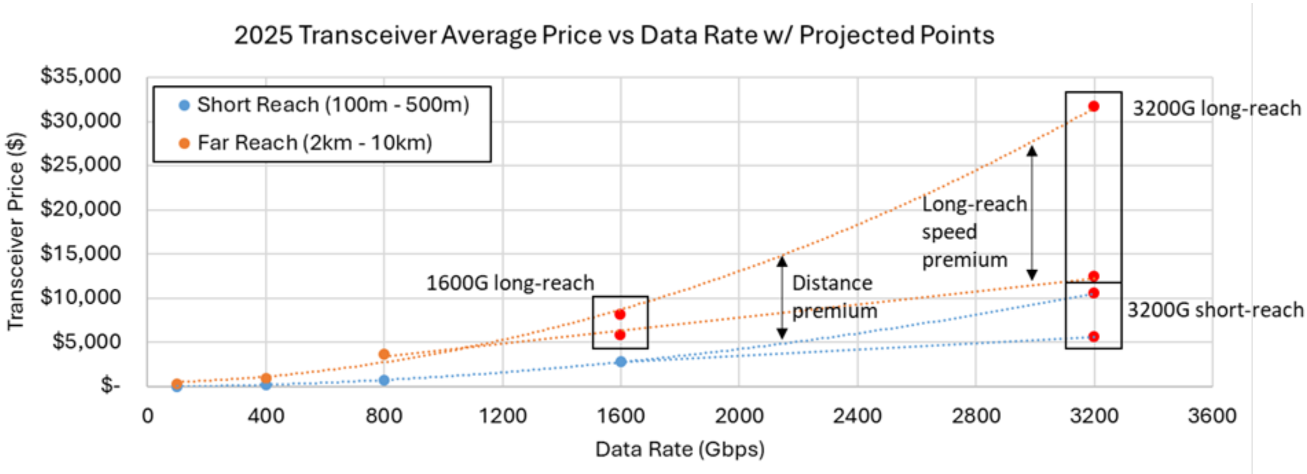

Figure 11. *Projected price points for products that do not exist in the market today are shown in red.*

The efficiency advantage may enable extended ranges provided the same transceiver input power. At a laser diode efficiency of 90% versus 40% for DFB lasers, a thermal management system would not draw as much power, and this excess power could be devoted to lasing. Conservatively assuming 2x range extension, a 3.2T long-reach transceiver could command prices in today's market from $12,000 to $31,000 – between 4x and 12x the price of current 500m range 1.6T transceivers.

Figure 11 shows both cases. The price range reflects two scenarios: the lower bound assumes that buyers become price-sensitive and premium payment patterns break down as transceivers reach extreme performance levels, while the upper bound assumes that the exponential speed premiums observed from 100G to 1.6T continue to 3.2T and beyond.

These dynamics reinforce the importance of the two breakthrough technology attributes discussed earlier. OAM directly addresses the strong price premium buyers place on speed. Faster diode modulation coupled with efficiency gains, meanwhile, offers the possibility of capturing value in both premium dimensions by delivering higher data density and extended reach. Together, these advances represent the type of step change the market is signaling demand for, rather than incremental brute-force scaling.

Product Mix and Prices over Time

Figure 12 projects baseline market demand curves, based on transceiver data rate only, through 2040 by building on the historic information presented in [14]. The projections assume transceiver data rate doubles every five years, creating a new product generation with a 20-year demand lifespan.

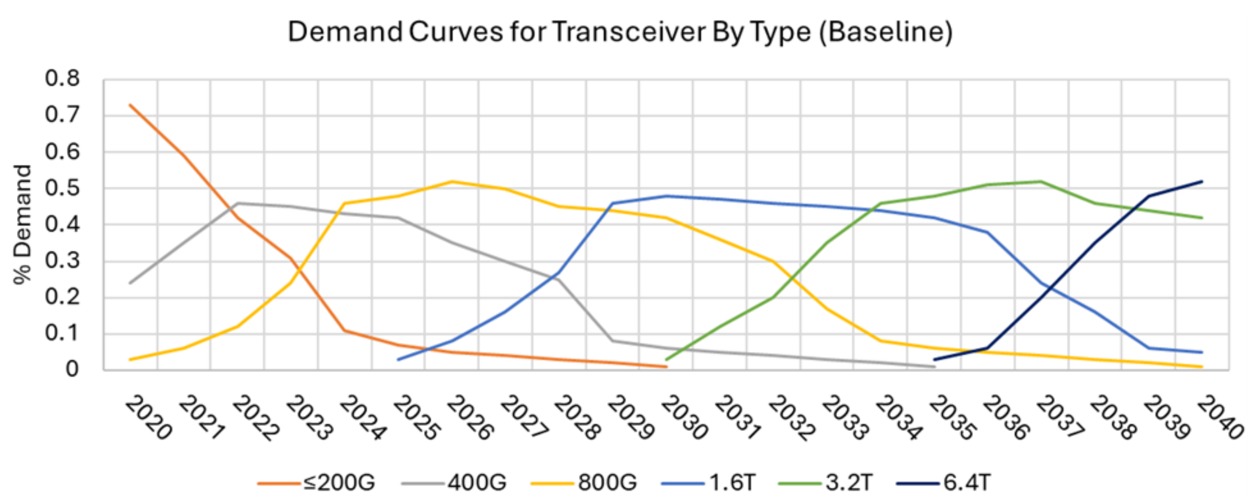

Figure 12. *Baseline Product Demand Curves.*

New technology could enable a generational leapfrog that pulls the product development lifecycle four years to the left. This would bring 3.2T transceivers to the market in 2026 instead of 2030, and it brings subsequent generations to market four years earlier. This trend leads to the cannibalization of 1.6T demand, as shown in Figure 13. Going forward, the term "leapfrog" is used to describe this scenario.

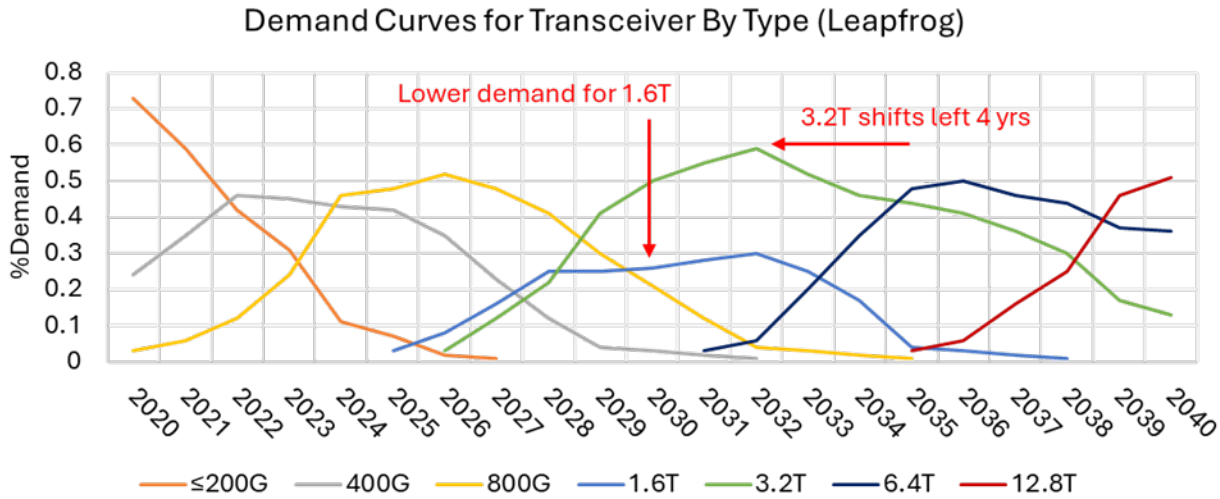

Figure 13. *Technology Leapfrog pulls the 3.2T product generation and all subsequent generations to the left by four years.*

Transceiver pricing changes with respect to demand, which follows new product generations. This analysis assumes the pricing for a product generation follows an exponentially decaying function with respect to time where each family has a value in the year 2025, established by the premium and no-premium trends presented in Figure 11. Figure 14 shows an example of short-reach transceivers subject to premium pricing, built off the baseline demand signal (Figure 12). Sixteen pricing curves similar to Figure 14 are used to define the price of short-reach and long-reach transceivers, each subjected to premium and no-premium pricing for the baseline and technology improvement scenarios described in the next section.

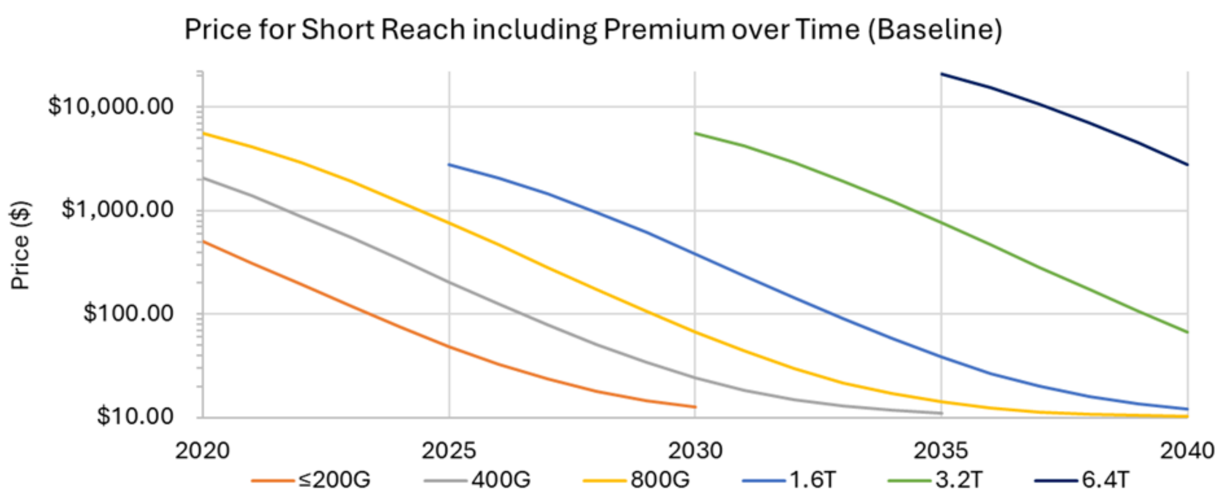

Figure 14. *Price of a short-reach transceiver subject to premium pricing. The first data point on a line is the first year where the demand is non-zero.*

Using the demand forecast, pricing estimates, and the assumption that the ratio of short-reach to long-reach transceivers is 80:20, enables projection of average transceiver price for the baseline and leapfrog cases broken into premium and no-premium. Four of the eight curves are shown in Figure 16 to provide an example.

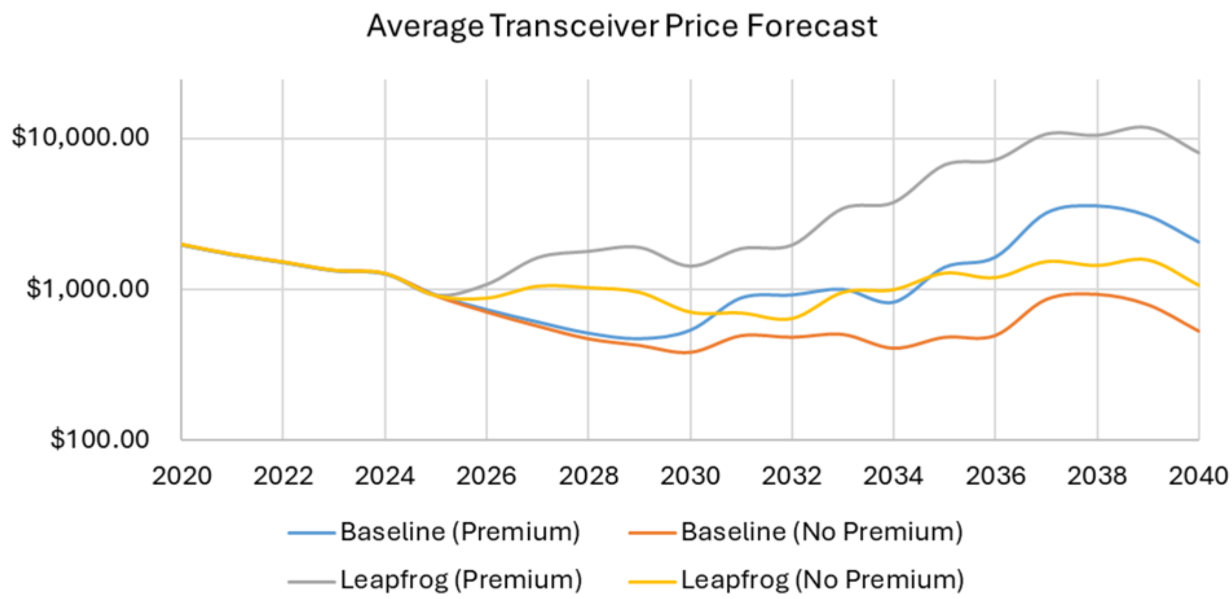

Figure 15. *Average transceiver price forecast based on a role-up of demand mix and price for the baseline and leapfrog scenarios, including premium and no-premium pricing.*

Projecting the Value of New Diode Technology

A sales baseline for optical transceivers is established through 2040, accounting for both replacement of failed units and natural market growth. Market growth estimates range from a Compound Annual Growth Rate (CAGR) of 13%[11] to 30%[14] through 2030. For conservatism, the baseline uses 13% CAGR but extends it through 2040.

Modeling assumptions include:

- **Geographic scope:** US market
- **Current Demand:** 14MM units of various data rates sold in 2025[14]
- **Market growth:** Growing 13% annually through 2040[11]
- **Failure model:** Weibull distribution with a 10% failure probability at five years due to a thermal degradation wear-out mechanism

The unit sales baseline shown in Figure 16 establishes a trend to provide the analytical foundation for quantifying how technological improvements translate to measurable market value in subsequent analysis. While the natural market growth rate is estimated at 13% CAGR, the bottom-up demand and pricing forecast methodology restrict revenue growth to 9% in the baseline. The baseline analysis predicts sales through 2040 **generate between $563 billion and $1.75 trillion in revenue.**

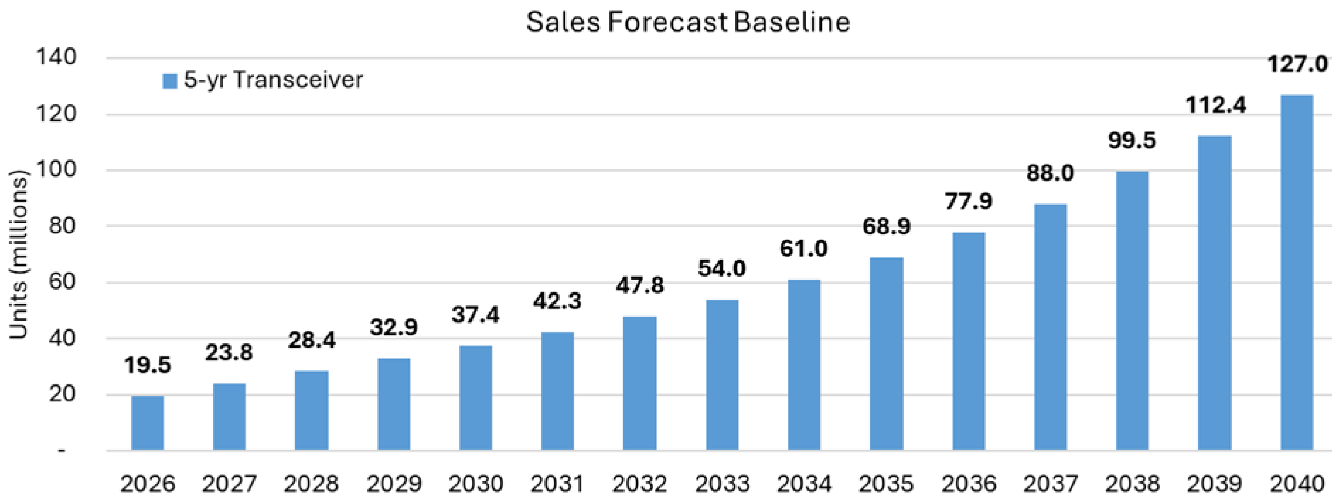

Figure 16. *Sales forecast of optical transceivers through 2040 include sales due to natural market growth and replacements for failures.*

There are three scenarios that could result from the introduction of new technology in 2026:

1. **Life Increase Only**: Diode efficiency improvements increase transceiver life from five years to ten years, doubling the sales price of generations beyond 1.6T.

   Improving transceiver efficiency significantly reduces heat, which could theoretically extend the laser diode's lifespan dramatically. The 10x lifespan estimate derives from reliability models showing that the heat reduction achieved at 90% efficiency (versus 40%) decreases operating temperature by approximately 10°C, which exponentially extends component life. However, such a dramatic increase in laser diode reliability would likely reveal other component weaknesses, particularly electrostatic discharge (ESD) failure modes in supporting electronics. While ESD failures are becoming less common as transceiver technology improves, they would likely become the limiting factor for overall system reliability. A more realistic overall transceiver life extension is 2x rather than 10x.

   Improved reliability could enable new data center architectural approaches, such as reduced redundancy requirements, though these secondary benefits are not quantified in this analysis.

   This analysis assumes current five-year transceivers are phased out with new models designed to last 10 years. For the calculations, the "end of useful life" (or 'life-limit') is defined as the point when 10% of the transceivers have failed. This prediction relies on a Weibull function that indicates a strong wear-out pattern, meaning failures become much more likely and rapid as the devices age (specifically, with a shape parameter of five).

   As shown in Figure 17, applying the improved reliability model to the sales projection indicates a reduction in total transceiver sales year-over-year as more reliable transceivers enter the market. However, this does not mean revenue is reduced.

   The pricing model for reliability improvements is based on the linear relationship established in Figure 9, which shows that users are willing to pay proportionally for increased data transmission over a transceiver's lifetime, but do not pay premiums beyond this direct relationship. While extended transceiver life reduces installation and replacement labor costs, inventory carrying costs, network downtime risks, and supply chain dependencies, market evidence suggests these operational benefits do not command pricing premiums beyond the direct cost-per-data relationship.

Therefore, doubling the life of the transceiver without any other performance improvements could double its price, but not more. This pricing strategy is applied to the 1.6T generation of transceivers and beyond beginning in 2026 and propagated through 2040 with the "average transceiver price forecasting" method.

The life increase generates between $895Bn and $2.86Tr in revenue between 2026 and 2040, representing a 1.6x increase over the baseline projection.

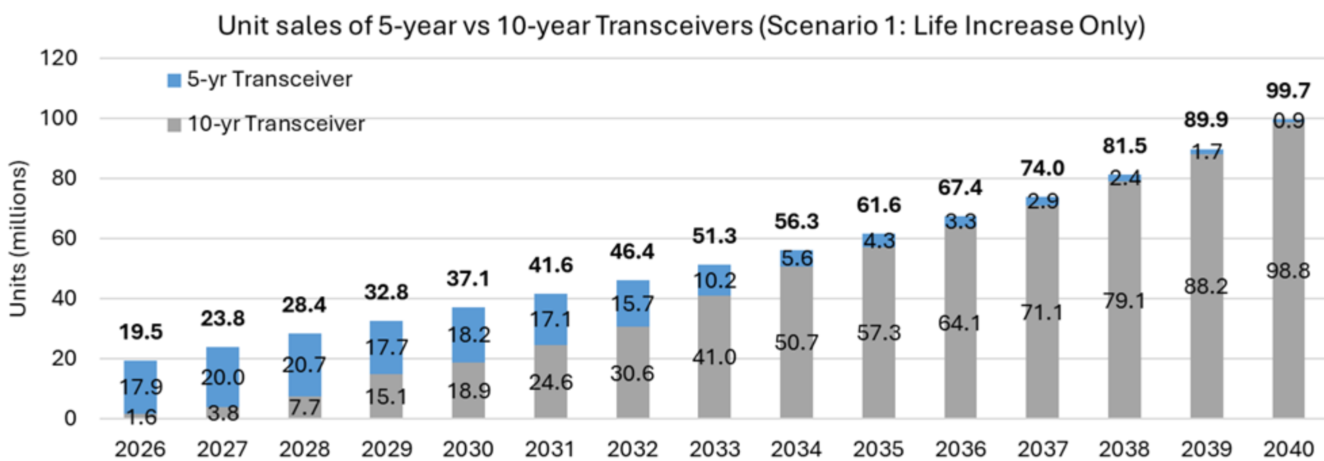

Figure 17. *10-year transceivers lead to a reduction in sales compared to the baseline.*

2. **Leapfrog, No Life Increase:** Diode efficiency improvements coupled with faster modulation enable 3.2T transceivers to enter the market in 2026, but without a transceiver life increase.

   This scenario focuses on sending data faster and farther but assumes no reliability increase because it is possible the fundamental failure mode remains when switching to new diode laser technology. Figure 13 shows the demand forecast used for this scenario and Figure 15 shows the average transceiver price forecast. In this scenario, 3.2T transceivers cannibalize the 1.6T sales while selling for substantial premiums, and future generations arrive four years prior to the baseline demand projections.

   The technology leapfrog generates between $1.07Tr and $6.22Tr in revenue between 2026 and 2040, representing between a 1.9x and 3.6x increase over the baseline projection.

3. **Leapfrog, With Life Increase:** Diode efficiency improvements coupled with faster modulation enable 3.2T transceivers to enter the market in 2026 with a transceiver life increase, doubling the sales price of generations beyond 1.6T.

This scenario combines the first two to produce the best possible value generating outcome. The reliability improvement is applied in the same manner as scenario one, and the technology leapfrog is applied in the same manner as Scenario 2.

The technology leapfrog with a life increase generates between $1.71Tr and $9.83Tr in revenue between 2026 and 2040, representing between a 3.0x and 5.6x increase over the baseline projection.

Figure 18 shows the results of all three scenarios relative to the baseline case in a single chart. Revenue multiples are reported above as the ratio of scenario-to-baseline results for each pricing case, showing the relative uplift from technology improvements.

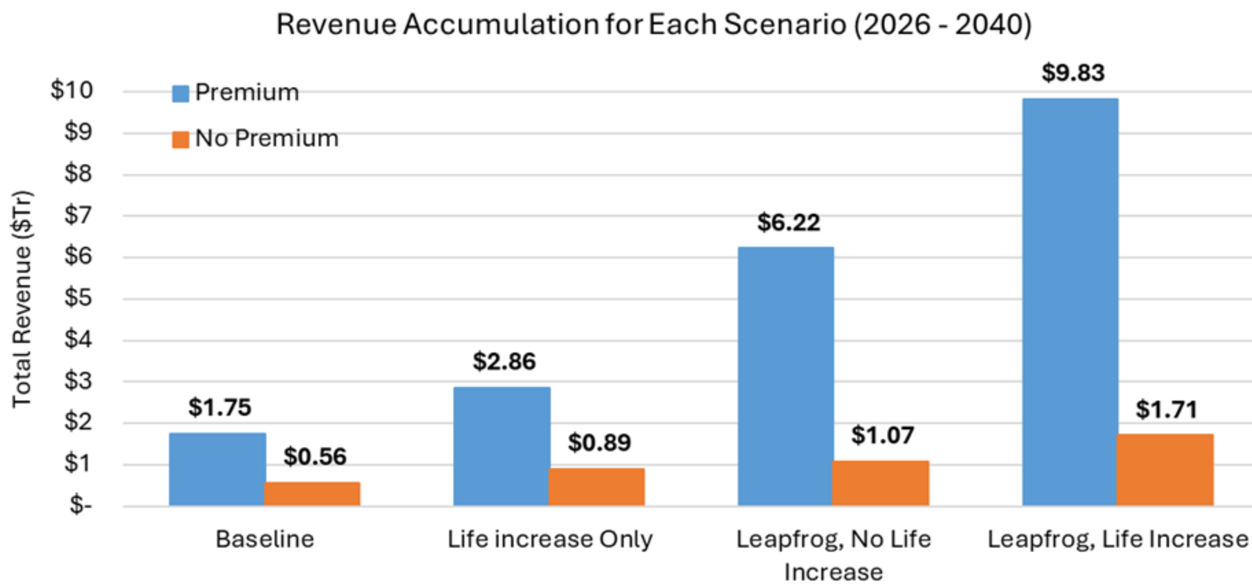

Figure 18. *Revenue accumulation for each scenario through 2040.*

While $9.83Tr represents the high-end scenario, the AI and edge infrastructure buildout could justify such sales projections. The global AI market is projected to exceed $1Tr annually by 2030, with data center infrastructure representing a substantial portion of this spending. Similarly, edge infrastructure deployment requires massive optical connectivity investments. When viewed against these macro trends, premium pricing for breakthrough transceiver technology that enables next-generation AI capabilities and edge performance becomes economically rational for competitive advantage.

Spending Less for Electricity:

Despite efforts to reduce transceiver power consumption, increasing the efficiency of a laser diode to reduce the electrical cost of a data center has minimal return. This is because compared to the other components within the transceiver, laser diodes already consume a minimal amount of power. Most power is devoted to the other driving electronics. Taking, for example, the 800G transceiver in Figure 3 (current 800G transceiver) and replacing it with the transceiver in Figure 6 (Scenario A from the engineering analysis) results in a 0.1W reduction in the power required to operate the transceiver, including associated cooling requirements.

The average industrial power rate in Northern Virginia, which has the largest concentration of data centers in the world, is 8.85 cents per kWh[8]. The energy cost savings of 0.1W, therefore, is minimal at $0.04 per transceiver per year.

Assuming a hyperscale data center has 100,000 800G optical transceivers, the cost savings for a single hyperscaler comes out to $4,000 per year, and across all 500 hyperscale data centers in the United States, it saves $2 million annually.

Even accounting for the transition to higher data rate transceivers (1.6T, 3.2T), the absolute power savings per transceiver remain small relative to total system power consumption. The efficiency gains continue to represent a small fraction of overall transceiver power requirements, making electrical cost savings negligible compared to transceiver acquisition costs of hundreds to tens-of-thousands of dollars per unit.

While electrical savings alone do not justify the technology investment, they provide additional value when combined with the performance and reliability benefits discussed previously. The primary value proposition for improved laser diode efficiency lies in enabling higher performance and longer lifespans rather than reducing operational power costs.

Final Remarks

The convergence of AI proliferation and edge deployment has created critical demand for optical transceiver technology. Incrementally increasing power consumption with each generation to achieve higher data rates is unsustainable as data centers approach power grid limitations. Advanced laser diode technologies utilizing Orbital Angular Momentum and designs that pulse faster with improved efficiency offer a paradigm shift that breaks this brute force trajectory.

The market opportunity is substantial and multi-faceted. Performance improvements enabling faster data transmission and extended range represent trillions of dollars of addressable market between 2026 and 2040. Reliability enhancements that double transceiver lifespans represent between hundreds of billions to a trillion in value in addition to those performance improvements. While electrical cost savings prove minimal in isolation, they provide supplementary benefits when combined with performance and reliability gains.

With hyperscalers aggressively competing for AI dominance and telecommunications providers racing to deploy edge infrastructure, the demand for breakthrough optical connectivity solutions is acute. For private investors, this represents a rare convergence of massive market opportunity, established demand patterns, and technology readiness. The near-term market entry of advanced laser diode technology requires aggressive investment and scaling, but it aligns with critical infrastructure buildout for AI data centers. Early investment in companies developing these technologies position investors at the forefront of a market transformation that could fundamentally reshape the digital infrastructure landscape.

The path from laboratory breakthrough to commercial success requires sustained private investment and rapid scaling. The trillion-dollar opportunity is real, measurable, and achievable.

BCS Allegient

Sources

[1] Mahdavi, M., Narayanan, V. "AI Driven Innovation for Directed Energy Systems." Pennsylvania State University (PSU) Center of Excellence in Directed Energy (CEDE) Kick-Off Meeting. April 23, 2025

[2] P. Ossieur, B. Moeneclaey, G. Coudyzer, J. Lambrecht, J. Craninckx, E. Martien. et. al. "Integrated Photonics and Electronics for Optical Transceivers Supporting AI/ML Applications." IEEE Journal of Selected Topics in Quantum Electronics, Vol 31. No. 3, May/June 2025. DOI: 10.1109/JSTQE.2025.3539379

[3] B. Srivathsan, M. Sorel, A. Bhan, P. Sachdeva, H. Batra, R. Sharma. el. al. "AI power: Expanding data center capacity to meet growing demand." McKinsey & Company. Octboer 29, 2024. Link: https://www.mckinsey.com/

[4] Synergy Research Group. "Hyperscale Data Center Count Hits 1,136; Average Size Increases; US Accounts for 54% of Total Capacity." March 19, 2025. Link: https://www.srgresearch.com/

[5] Lumentum. "3.2W 940nm VCSEL Array." Datasheet. 2022.

[6] C. Minkenberg, R. Krishnaswamy, A. Zilkie, D. Nelson. "Co-packaged datacenter optics: Opportunities and challenges." *IET Optoelectron*. 2021;15:77–91. DOI: 10.1049/ote2.12020

[7] L. Li, Y. Xiao, W. Wang, C. Guan, W. Yao, Y. Zhang. el al. "Advanced Fabrication of 56 Gbaud Electro-Absorption Modulated Laser (EML) Chips Integrated with High-Speed Silicon Photonic Substrates." *Photonics 2025*, 12, 329. Link: https://doi.org/10.3390 /photonics12040329

[8] U.S. Energy Information Administration (EIA). "Electric Power Monthly. Average Price of Electricity to Ultimate Customers by End-Use Sector." Link: Electric Power Monthly - U.S. Energy Information Administration (EIA)

[9] R. Perez-Aranda, D. Ortiz. "VCSEL reliability comparison." *IEEE 802.3cz Task Force*. June 8, 2021.

[10] Y. Sin, J. Theiss, M. Huang, E. Tang, J. Childs, A. Bushmaker, et al. "Reliability and Non-destructive Failure Investigate of 850-nm Multimode and 795-nm Single-Mode VCSELs for Space Applications." *Proc. of SPIE Vol. 13384, 1338409*. March 19, 2025

[11] MarketsandMarkets. "Optical Transceiver Market Size, Share & Growth Analysis." Report Code SE 4626. March 2024.

[12] CablesPlusUSA.com. "How Much Does Fiber Optic Cable Cost?" December 22, 2023. https://store.cablesplususa.com/cabling-insider-blog/fiber-optic-cable-cost/

[13] Coherent Corp. "Coherent Unveils High-Efficiency Lasers For Silicon Photonics Transceivers." August 28, 2024. https://www.coherent.com/news/press-releases/high-efficiency-lasers-for-silicon-photonics-transceivers

[14] B. Wiseman, H. Marcil, L. Hamalainen, P. Sachdeva. "Opportunities in networking optics: Boosting supply for data centers." June 2025. McKinsey & Company.

BCS Allegient