WHITE PAPER

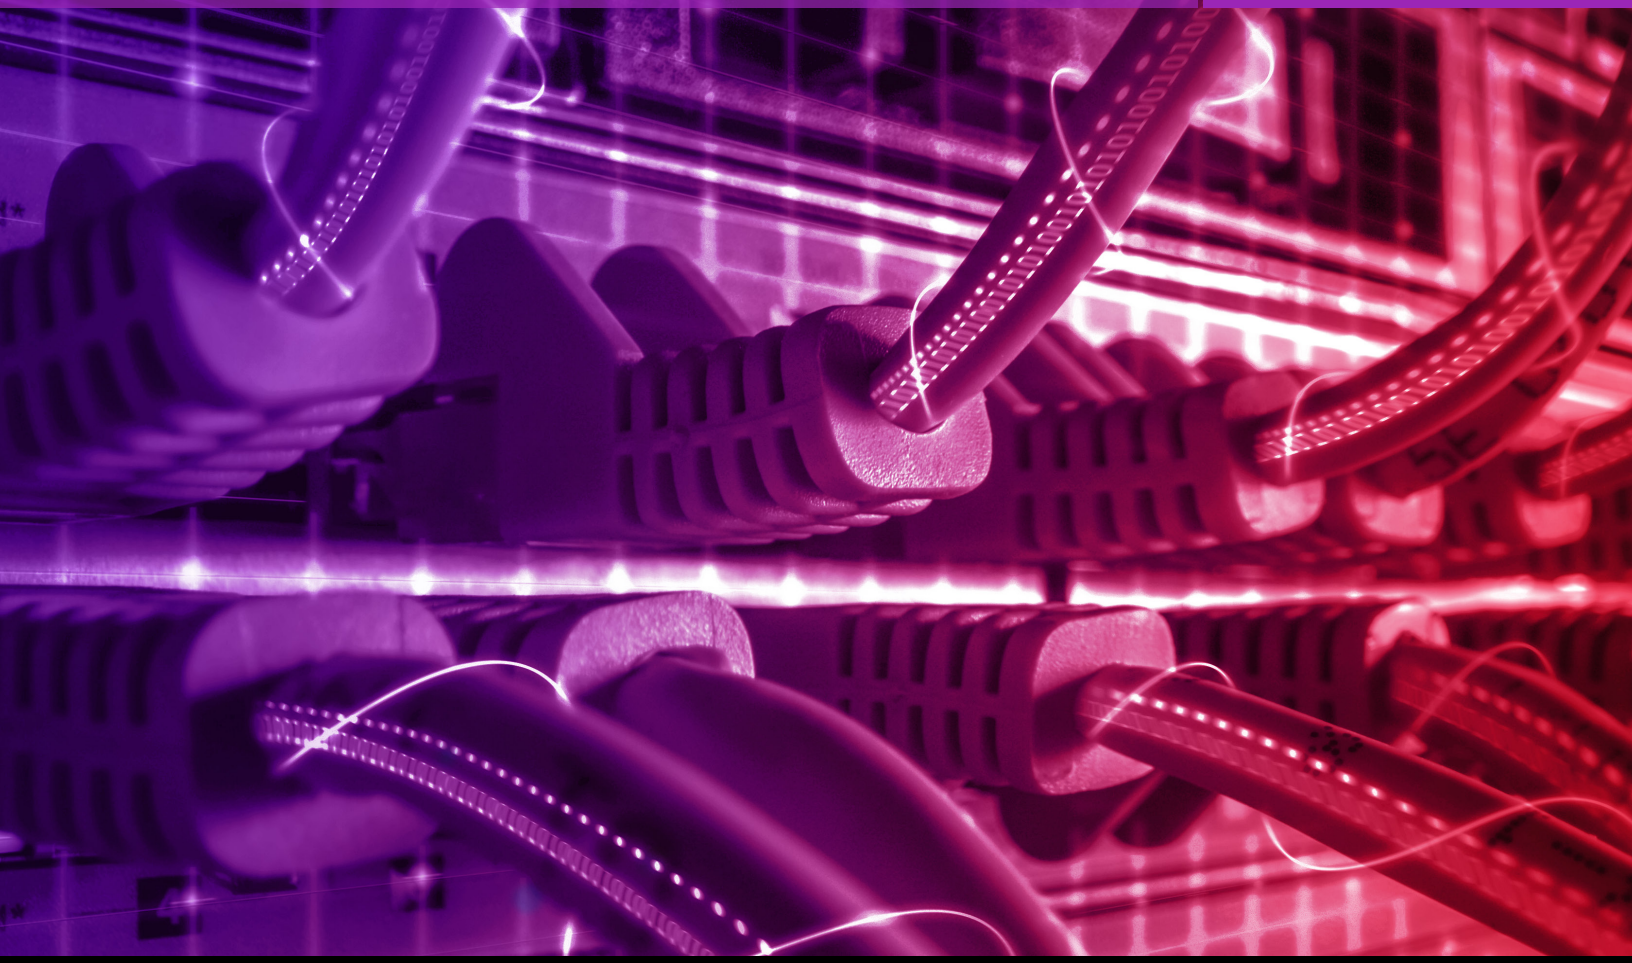

800G / 1.6T Data Center Transceiver Test

Overcoming three challenges to unlock AI data center potential

KEYSIGHT

Introduction

Artificial intelligence (AI) generates vast amounts of data, driving new computing, storage, and performance demands on data centers. AI clusters, servers containing hundreds of graphic processing units (GPUs), will become the standard unit for training machine learning (ML) algorithms. Data center operators need to embrace new architectures and technologies to support the response times and high bandwidth these technologies require. AI networks will look entirely different from today's data centers.

Today's 400 gigabit (400G) Ethernet networks will not be fast or power-efficient enough to manage machine learning algorithm training. Hyperscale data center architects and operators must scale networks from 400G to 1.6 terabit (1.6T) Ethernet to meet these demands.

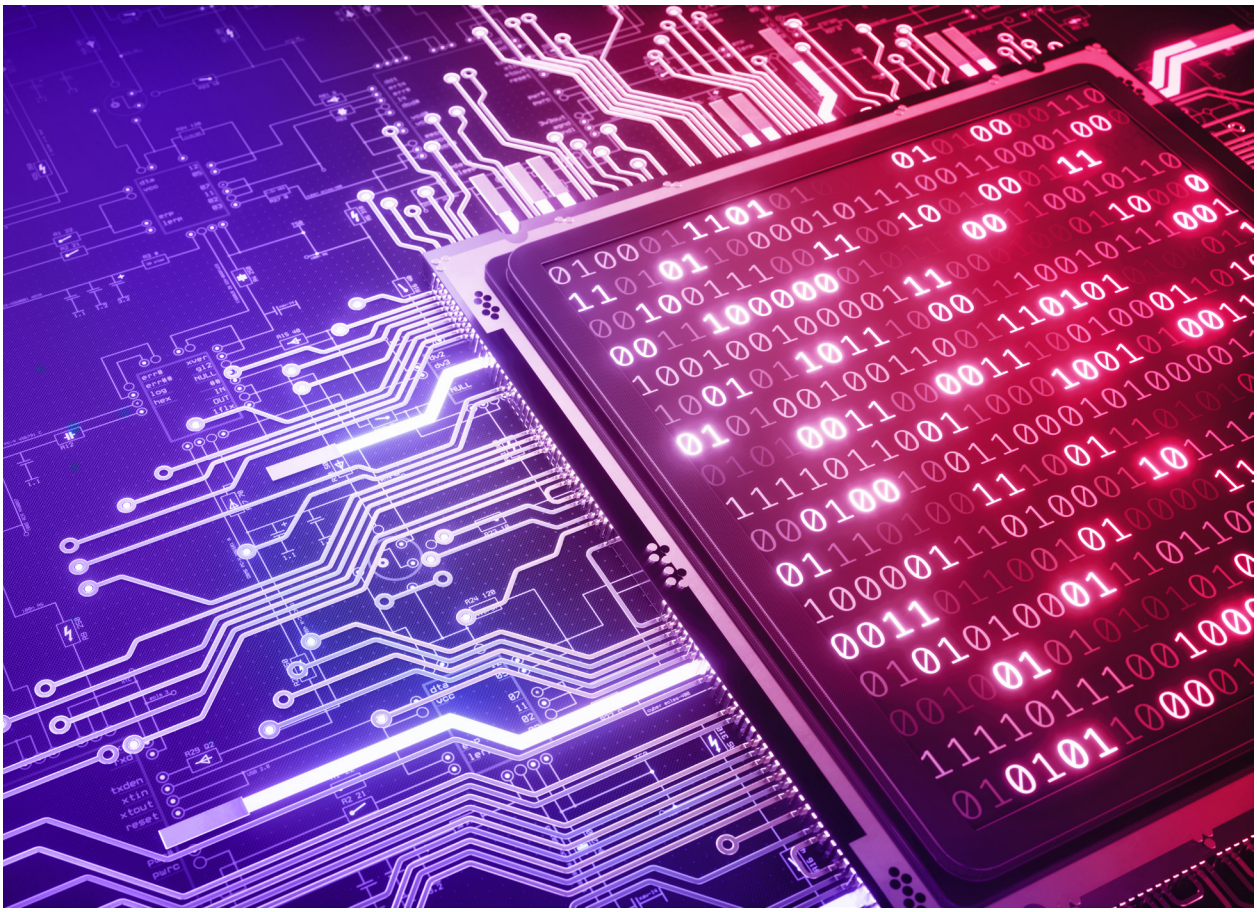

AI Data Center Trends

Machine learning training requires data centers to scale out by adding more interconnects between compute units. As illustrated in Figure 1, East-west traffic is projected to increase significantly to support AI training, driving transceiver growth from thousands to millions of interconnects per data center in the coming years. Each of these interconnects will have to increase in data speed as well, increasing network speeds to 1.6T. Every connection is a potential weak link. One failed or unoptimized transceiver could disrupt an entire AI workload, wasting considerable time and money.

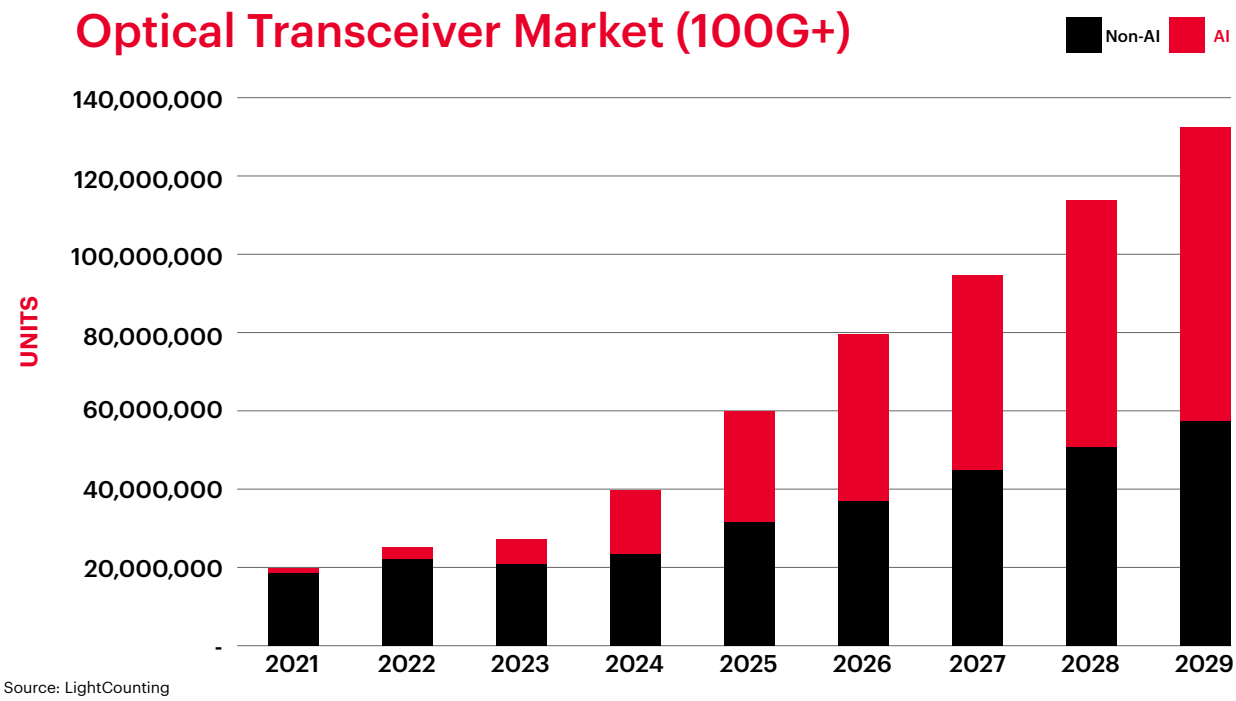

Figure 1. AI is driving growth in the data center transceiver market (*Source: LightCounting Market Research*)

The transition from 400G to 1.6T

Data center speeds have swiftly evolved over the past two decades from 10G to 800G. Each new speed class was built upon the previous generation but faced limits of high cost and constrained data capacity. Traditionally, service providers were the first to test transceivers for the next speed class in the network context. However, because of increased networking bandwidth and reliability needs from each component during AI / ML training, network equipment manufacturers, especially optical transceiver manufacturers, must ensure high-quality devices through rigorous testing at both the physical layer and protocol / network layers.

Furthermore, data center operators face different requirements and challenges than service providers. For example, modern hyperscale data centers house more than 50,000 fibers with an optical transceiver at each end, so transceiver cost drives data center operators to pioneer technology. The following are three key challenges they need to tackle to ensure that data centers support 800G and 1.6T speeds:

- Increase interconnect bandwidth.
- Ensure interoperability and reliability.
- Reduce test time, power, and cost.

Challenge 1: Increase interconnect bandwidth

Data center operators traditionally upgrade their network architecture every few years. However, many data centers are at maximum capacity, and operators need to find a way to process more data faster to reach 1.6T speeds. Data center operators rely on transceiver manufacturers to move to the next generation of speed class.

Solution: Faster switching silicon and forward error correction

Advancement from non-return-to-zero (NRZ) signaling to four-level pulse amplitude modulation (PAM4) enabled speeds to jump from 100G to 400G. Now, 56 Gb/s optical and electrical data rates were possible at 28 GBd/s symbol rates. Soon after, 112 Gb/s lanes followed. To reach 1.6T speeds and satisfy the demand for faster data centers, developers need 224 Gb/s lanes. As seen in Figure 2, the time between generations is compressing as the industry tries to keep up with demand. Developing devices capable of communicating via 224 Gb/s optical and electrical lanes requires higher bandwidth oscilloscopes and bit error rate testers with enough signal integrity to properly characterize transmitter and receiver performance with low noise and jitter margin.

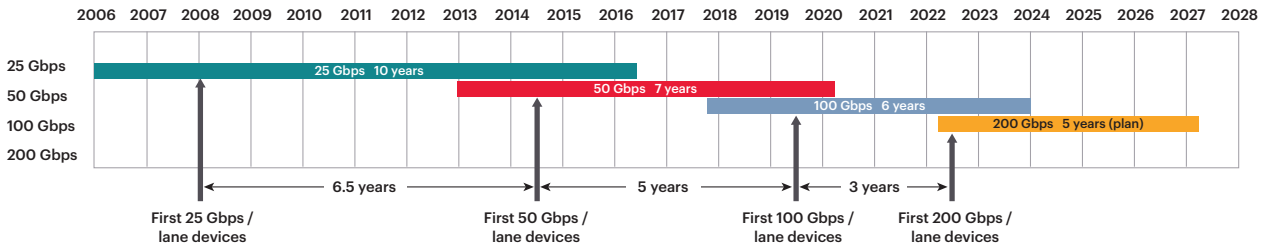

Figure 2. Timeline of interconnect data rate advancements over time

The road to faster lane speeds has many components. Complementary improvements in metal-oxide-semiconductor (CMOS) process technology enable faster and more efficient switching of silicon. Innovations in high-bandwidth-switching silicon from 25.6T to 51.2T enable faster lane rates without increasing server switch radix. Digital signals become more susceptible to channel loss at these faster data rates. One solution is implementing more sophisticated forward error correction (FEC) algorithms to minimize symbol loss. FEC is an advanced coding technique that sends the information to correct errors through the link along with the payload data. Testing at Layers 2 and 3 to validate operation at the protocol and network layers using an interconnect and network tester is essential for reducing unexpected bit errors from a build-up across multiple lanes.

Challenge 2: Ensure interoperability and reliability

Before being inserted into the network, any new transceiver technology used in pluggable modules requires thorough testing to ensure that it complies with specifications to offer seamless compatibility. Optical transceiver manufacturers must test to ensure that their transceivers meet defined specifications and interoperate with other network components and transceivers from different vendors. For network operators, network downtime due to faulty transceivers is unacceptable because of service-level agreements. In AI data center networks, it could be even more costly. Developers must identify performance issues before their components become weak links in the data center, causing network shutdowns and costly workload failures for their customers, the data center operators.

Solution: Characterization and compliance test

Several standards organizations govern optical transceiver specifications and define test procedures to ensure compliance with standards and interoperability with other vendors. They include the Institute of Electrical and Electronics Engineers (IEEE), the International Committee for Information Technology Standards, and the Optical Internetworking Forum. Data center operators can ensure the quality and compatibility of their next-generation data centers by selecting a transceiver vendor that has successfully passed all physical Layer 1 characterization and compliance tests defined by industry standards.

Standards organizations define specifications and provide compliance test procedures to ensure that a receiver will operate with a worst-case transmitter and vice versa. Transmitters and receivers need different sets of optical and electrical tests. The effects of the channel between them also require consideration. Faster and more complex systems increase the difficulty and time of test. High bandwidth, high precision oscilloscopes, high-performance bit error ratio testers, and test automation software can help developers ensure that their devices not only pass the increasingly complex test requirements but exceed compliance thresholds under real-world conditions.

Challenge 3: Reduce test time, cost, and power

The cost and time commitment to test next-generation optical transceivers is a major hurdle for data center operators transitioning to 1.6T. The cost of transceivers is directly proportional to the complexity of the design and the number of optical components. Test time is another factor in the cost.

The complexity of the test systems will also increase as devices under test will have 8, 16, or 32 lanes. The need to measure complex specifications according to standards complicates the design and validation process and requires a long learning curve for test engineers. The move to PAM4 reduced the margin for error in test processes because of its tighter thresholds. This margin was further reduced in 224 Gb/s lanes.

Cost and power consumption define the overall footprint of a transceiver. Hyperscale data centers consume about three percent of the power generated worldwide. Data center technology needs to reduce power consumption per bit and cost per bit to maintain the current growth rate and scale with the need for higher bandwidth.

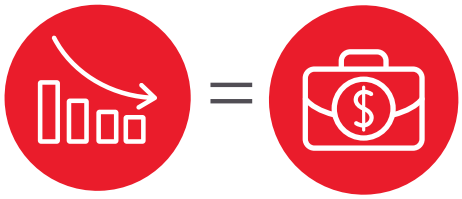

Solution: Test and power efficiency

The proliferation of communication standards and transceiver types during the last decade has created more complexity for R&D and manufacturing test. Fortunately, several techniques significantly reduce test time and the overall cost of transceivers, from design and simulation through device characterization and compliance to, finally, manufacturing. Tighter test margins require more precise, higher-bandwidth instruments.

Automated compliance testing eases the learning curve, assists with debugging, and increases test efficiency by automating test procedures. Optical test optimization software helps quality assurance (QA) engineers improve hardware utilization of fixed configuration sampling oscilloscopes to improve the efficiency of manufacturing testing. High-throughput, high-accuracy transmitter dispersion and eye closure quaternary (TDECQ) testing will be essential for scaling 1.6T transceiver manufacturing to prepare for the growth of interconnect demand from AI data centers.

New technologies can help increase power efficiency. The faster, more efficient switch silicon and 224 Gb/s lane rates help process twice as much data while lowering power consumption. Some technologies will require test time, cost, and power efficiency trade-offs. For example, linear pluggable optics and co-packaged optical modules lower the power consumption of the module per bit compared with pluggable optics by removing the digital signal processing (DSP) from the transceiver itself and relying on the host processor for equalization. This new topology creates additional challenges when testing with proper equalization. Similarly, photonic integration, enabled by silicon photonics, helps drive down the cost of assembly but requires more advanced testing.

Summary

The AI revolution is here, and 1.6T optical transceivers represent the next step in network bandwidth for AI-capable data centers. Transceiver manufacturers that increase bandwidth, ensure reliability, and optimize test time, cost, and power consumption will lead the transition to 1.6T Ethernet powering the next generation of AI.

Explore how Keysight solutions address the challenges of 1.6T data center and AI infrastructure:

- Learn about Keysight AI solutions to scale data centers from the physical to the application layer.
- Hear about the latest in 1.6T development from data center industry experts.
- Read more about the challenges of developing and testing at 224 Gb/s, Data Center Ethernet Technology and Evolution to 224 Gbps – Application Note.