

400ZR: Accessible 400G for Edge DCIs and Beyond

Developing a practical 400G coherent transceiver in a compact form factor requires innovation and integration.

Eugene Park, Ph.D.

Rapidly growing content demand is pressuring carriers, cloud providers, and traditional data center operators to boost the bandwidth of the data center interconnects (DCIs) that link their facilities. A number of factors contribute to the situation. By the year 2022, roughly 82% of all IP traffic is forecasted to be in the form of streaming video.¹ An estimated 4.8 billion global Internet users are forecasted to be accessing content from an Internet comprised of 28.5 billion networked mobile devices and connections. Meanwhile, emerging applications such as the Internet of Things (IoT), 5G wireless, and ultra-high-density (8K) video are poised to intensify demand. Network operators are planning changes to their architectures to keep pace with this growing demand for content.

To address the need for higher DCI bandwidth requirements to meet this growing demand, the optical networking industry began working on a solution known as the 400ZR implementation agreement, with a goal to combine optical line-side fiber capacity with the benefits of client-side solutions.

Spearheaded by the Optical Internetworking Forum (OIF), 400ZR aims to deliver accessible 400 gigabit-per-second (Gbps) Ethernet links for edge DCI applications (see figure 1). The 400ZR implementation agreement addresses edge-DCI applications with link distances targeting 80 km to 120 km and can be implemented in pluggable 400Gbps optical transceiver module form-factors used for client optics.

The 400ZR standard leverages the advantages of coherent optical technology, in which optical laser light is manipulated using a combination of phase and amplitude modulation to transmit multiple bits per symbol (see *Coherent Optics Basics* sidebar). Optical coherent technology capable of 400Gbps operation is already deployed in carrier networks. The challenge is to transform this technology to meet the requirements of 400ZR and implement the technology in pluggable module form-factors used for client optics.

This market backgrounder reviews the industry drivers behind the development of 400ZR, the key benefits of the technology, and a product-development roadmap for bringing 400ZR transceivers to market. It also provides a few predictions on how the technology could potentially change the industry.

Practical high-bandwidth DCIs

Data centers began as off-site installations intended to support the primary business facility,

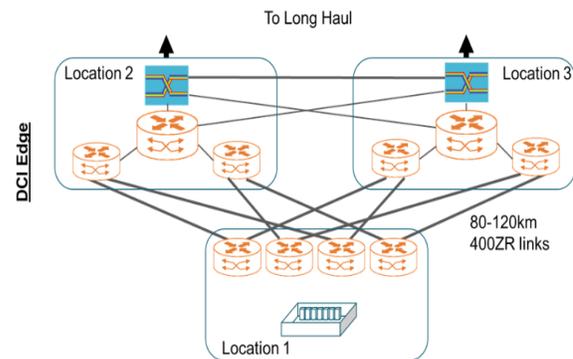


Figure 1: The primary 400ZR use case is to apply the technology to DCI edge networks.

whether for backup, disaster recovery, or because of space or power constraints. Today, data centers have become the primary business facilities for enterprises such as Internet content providers, cloud services providers, communications services providers, and more. Although data transmission to the end-user (north-south traffic) remains of primary importance, these new business models require a far greater amount of internal data transmission (east-west traffic). Enterprises want to be able to add additional facilities but manage the collection as a single asset. By linking inter-data center networking equipment over high-speed optical interconnects, data center operators plan to create scalable, resilient virtual hyperscale data centers.

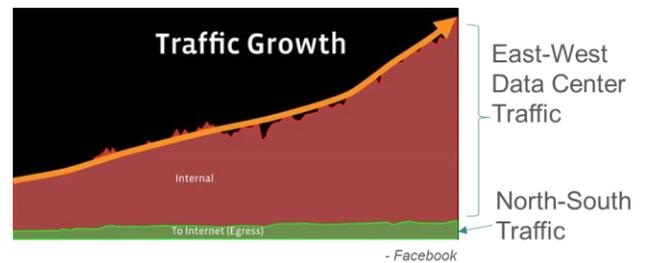


Figure 2: In today's hyperscale data networks, east-west traffic (inter- and intra-data center) dominates total traffic.

In addition, the most recent generation of Ethernet switching chips, which offer 12.8Tb operation and 400Gbps I/O (input/output) capabilities, should soon see broad deployment in these hyperscale data centers. These I/O ports depend on 400Gbps optical transceiver modules to connect these switches over the data center's fiber optic infrastructure. This is where 400ZR plays a serious role in the growth of these data centers.

Practical Solutions

Deployment of 12.8Tb switches with 400ZR optical transceivers in hyperscale data centers should address a number of data center pain points such as cost per bit and density. Let's discuss how 400ZR addresses these pain points.

Fast, Streamlined Architectures

Cost per bit is a key metric for data center operators. With the increased focus on

transmission capacity, the challenge becomes how to develop network architectures suitable for the data center with hardware and software solutions capable of delivering bandwidth at an acceptable cost per bit. Data centers primarily use the Ethernet protocol, which provides an economical approach for switching/routing packets of data, referred to as Internet Protocol (IP) packets.

Given the scale of the modern data center, bandwidth alone is not sufficient. Network architectures also need to minimize the installation and management burden. With respect to the optical portion of the network connecting multiple data center sites, there are generally two ways to transport the IP/Ethernet traffic over an optical infrastructure. A traditional method is to create an optical transport dense-wavelength-division-multiplexed (DWDM) network layer and then attach IP/Ethernet switch/router equipment to this network via short optical cables (see figure 3). For ≤ 100 Gbps links

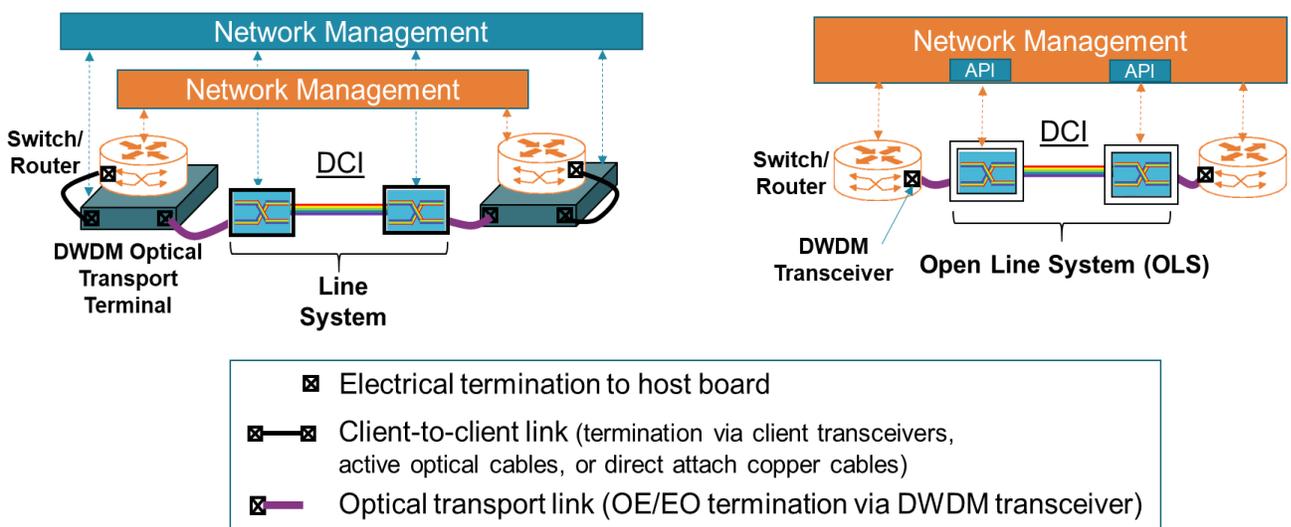


Figure 3: Two common architectures for transmitting IP/Ethernet traffic over an optical infrastructure are maintaining a separate DWDM optical transport network and IP/Ethernet transport network (left) or plugging DWDM transceiver modules directly into the switch/router boxes (right).

short copper cables may also be used. Using the optical transport terminal equipment, data from each switch/router port is “mapped” onto a distinct DWDM wavelength. The IP/Ethernet layer and the DWDM optical layer are treated as separately controlled networks.

Another implementation method is to create what is called an IP-over-DWDM (IPoDWDM) network in which optical transceiver modules with DWDM lasers are plugged directly into each IP/Ethernet switch/routers port. An IPoDWDM network has the potential of reducing both cost per bit and operational overhead since less equipment is required and network management can be consolidated.

Some network operators may find it necessary to maintain separate optical and IP networks, especially to support legacy infrastructure. Other operators may focus on the benefits of reducing the amount of equipment they have to manage, especially given scalability concerns.

400ZR optical transceiver modules enable data centers to deploy an IPoDWDM infrastructure by plugging 400ZR modules directly into 400Gbps Ethernet I/O ports on their switch/router platforms.

Familiar form factors

Space is typically a premium in the data center. Operators want to maximize the density, which refers to the number of data I/O ports that can fit within a rack. This translates to maximizing not only the amount of equipment that can fit within a rack, but also the number of optical module ports each rack-mounted terminal equipment can support. Data centers have widely adopted the multi-source agreement (MSA) defined quad small form factor pluggable (QSFP) family which provides high faceplate densities. From the 40Gbps QSFP+ variant to the 100Gbps



Figure 4: The two form factors under consideration for the 400ZR implementation agreement are the quad small form factor pluggable, double density (QSFP-DD, left) and octal small form factor pluggable (OSFP, right).

QSFP28 variant, the data rate supported has increased while the mechanical size has remained relatively unchanged. The latest variant, QSFP-DD, is under development to support 400Gbps.

The 400ZR implementation agreement for data center transceivers is intended to support QSFP-DD, as well as a competing form factor called octal small form-factor pluggable (OSFP, see figure 4). Most 12.8Tb switches and routers designs are expected to incorporate slots to accommodate QSFP-DD or OSFP modules. This is to enable 400ZR transceivers to be plugged directly into the switches to support IPoDWDM architectures. The central challenge for transceiver manufacturers is to find a way to implement coherent technology in these very compact packages.

Why coherent optical transmission?

Traditional optical transmission uses a form of intensity modulation called non-return to zero (NRZ) where the laser light is turned on and off at the bit rate, in other words, one bit per symbol. This NRZ signal is received using a technique called direct detection. Short distance 400Gbps pluggable transceivers are being developed using a form of intensity modulation with four-level amplitude transmission called

PAM4 (4-level pulse amplitude modulation), in conjunction with a direct detect receiver, resulting in two bits per symbol (see *Coherent Optics Basics* sidebar). These direct detect solutions rely on multiple lasers to achieve throughput as high as 400Gbps. However, they do not have the range required to achieve the distances needed for 400ZR. In contrast, coherent transmission sends multiple bits per symbol using a combination of phase and amplitude modulation, as well as transmitting in two orthogonal polarizations. For example, the 400ZR implementation agreement states the utilization of dual-polarization 16QAM (DP-16QAM) modulation with a baud rate of approximately 60Gbaud. This enables transmission at 400Gbps and beyond using a single laser source, while offering performance benefits that enable improvements in reach compared to direct detection.

Earlier generations of coherent modules tended to be significantly larger than their intensity-modulated counterparts because of the need for powerful digital signal processors and the use of discrete optical components. These solutions targeted long-haul applications where power was less sensitive. However, advances in CMOS, integrated optics, coherent digital signal processor (DSP) designs, as well as DSP coding and equalization algorithms, have significantly reduced the power and size of coherent interfaces for edge DCI applications while maintaining a high level of performance. Today, the implementation complexity between PAM4/direct-detect and DP-16QAM/coherent solutions are similar. Both solutions require integrated optics and custom DSPs, similar RF electronics and advanced packaging technology.

With rising data demand, data center operators are seeking a low power, small form-factor solution to cover edge-DCI distances. Hence,

the industry efforts to deliver a viable coherent solution—that viable solution is 400ZR.

The importance of interoperability

As previously stated, the 400ZR implementation agreement calls for a pluggable coherent 400Gbps solution based on DP-16QAM signaling with approximately 60Gbaud transmission rates. These devices allow for supporting 80 to 120 km links. A key advantage of standardization is the expectation of interoperability of 400ZR modules from different suppliers.

Historically, coherent optical networking equipment has been built around proprietary (closed) solutions. With 400ZR, coherent solutions more closely resemble MSA client optics with the expectation of interoperability such that a module from vendor A on one end can optically link up to a module from vendor B on the other end. This interoperability further supports an IPoDWDM infrastructure where switches are optically linked to each other over an open line system using 400ZR modules. An open line system built around 400ZR would give data center operators and carriers greater choice in their DCI network architectures and in the components they use to build them.

The roadmap to 400ZR

The expectation of the wide adoption of 400ZR by hyperscale network operators for edge applications is generating significant industry attention in 400ZR. Practical devices need to be designed and fabricated, and then produced in volume. Similar to the shorter reach

400Gbps Pluggable Variants

The optical coherent technology required for 400ZR can also benefit other markets. In response, two main variants have surfaced: 400ZR+ and Open ROADM 400Gbps. Open ROADM 400Gbps expands upon the 400ZR spec to include carrier-centric specifications, including higher performance, outlined by the AT&T-led Open ROADM MSA. 400ZR+ is functionally similar to 400ZR, with support for optical performance similar to Open ROADM. Many of the same suppliers supporting the 400ZR effort are also working on these enhanced-performance variants.

PAM4 solution, packaging optics into the QSFP-DD/OSFP form factors is challenging.

Complying with these mechanical designs while meeting requirements for performance, power consumption and cost requires an innovative focus on three key areas: the DSP, optical/electrical component consolidation, and high-density packaging.

The investment in 400ZR is also anticipated to benefit other applications in the broader optical networking community. For example, the technology developed for the 400ZR small-form-factor pluggable transceivers can also be applied to next-generation CFP2 modules.

Despite the march toward smaller form factors, the CFP2 is expected to remain widely used in carrier applications at 400Gbps. The extra space available in the CFP2 package can enable vendors to offer not just coherent transceiver technology, but also introduce additional features and functionality to better serve the carrier space. We refer to these pluggable modules targeting carrier applications as an Open ROADM 400Gbps variant (see *400Gbps Pluggable Variants* sidebar). Let's take a look

under the hood of a module to learn the challenges of manufacturing a 400Gbps transceiver in a QSFP-DD/OSFP form factor.

The DSP and Moore's Law

Complex optical modulation schemes and advanced error-correction algorithms require a high-functioning DSP. Thanks to Moore's law, which predicts transistor density doubling approximately every two years, a new generation of high-functionality, size-reduced DSPs can be leveraged for 400ZR and other 400Gbps coherent transceiver applications. These DSPs are key to achieving a compact size module with low power consumption.

Optical/Electrical Component Consolidation

Figure 5 illustrates how advances in optical/electrical component consolidation results in size reductions of coherent modules leading to QSFP-DD/OSFP. In addition to the size reduction of DSPs based on Moore's law, the use of silicon photonics to take discrete bulky optical components and integrate their functions into a CMOS-based silicon chip is largely responsible for the module footprint reduction. Let's take a look at how silicon not only provides a high level of integration, but also why silicon is an economical solution.

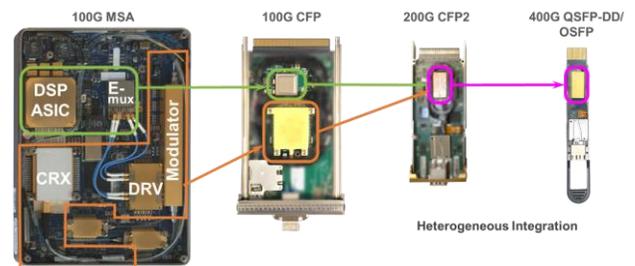


Figure 5: Heterogeneous integration in silicon enables a path to coherent transceivers in a QSFP-DD/OSFP package.

Materials systems

The high-speed optical components for coherent designs are generally based on either indium phosphide (InP) or silicon. InP is a direct bandgap semiconductor material well-suited for use in optical components such as lasers and modulators. Using InP, these functions can be fabricated on the same chip (see figure 6). In contrast, silicon is an indirect-bandgap semiconductor material, commonly used to fabricate various electronic components including DSPs and other ASICs. Silicon can also be used to create devices with optical functions (aka silicon photonics) such as modulators and coherent receivers. However, silicon does not allow for the straightforward fabrication of lasers. As a result, in a silicon photonics coherent design, the laser is typically separate from the rest of the optical circuit and made from a different material such as InP. This separation can be an advantage in thermally challenging environments because the temperature-sensitive InP laser can be designed to be further away from heat-producing ASICs, while the temperature insensitive silicon photonics can be placed closer to the heat source.

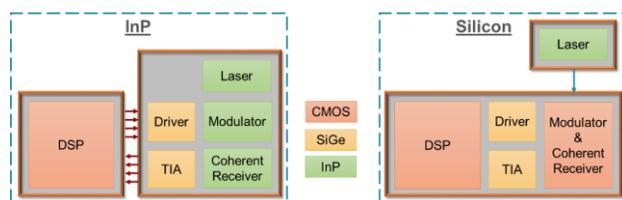


Figure 6: Indium phosphide (left) enables integration of everything but the DSP; and silicon (right), which enables the integration of many electronic components, but not the laser.

Whereas InP has limited application beyond optics, silicon photonics benefits from leveraging common CMOS processes. Transceiver manufacturers can draw on the resources of a

mature electronics industry, benefiting from volume manufacturing, lower costs, broad availability of fabs, and ongoing technology improvements.

Packaging Technology

An effective path to small-form-factor pluggable coherent transceivers is to follow the example of the electronics world and apply heterogeneous integration techniques such as 2D and 3D co-packaging. For example, electronic circuits and a silicon PIC can be packaged side-by-side on a common substrate or stacked (see figure 7). The heterogeneous integration approach makes it possible to integrate the PIC and other crucial components such as transimpedance amplifiers (TIAs), analog-to-digital convertors (ADCs), and digital-to-analog convertors (DACs) into a compact package. An advantage of this approach is the reduction of electrical interconnects while preserving robust signal integrity. Heterogeneous integration techniques are already used in the electronics industry for compute and memory chips. Heterogeneous integration involving silicon photonics follows progress already made in the semiconductor fabrication process suitable for volume production and high yields.

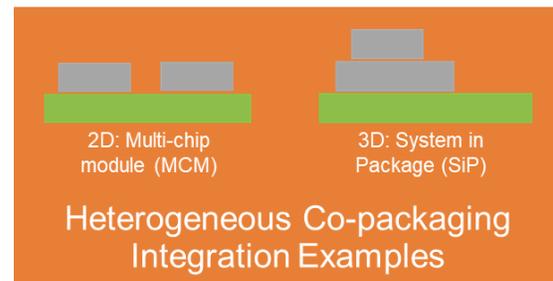


Figure 7: Heterogeneous co-packaging of silicon enables the integration of high-speed optics and electronics.

Meeting the 400ZR challenge

The 400ZR implementation agreement is a customer-focused effort to enable data center operators to meet increasing demand for content. Speed alone is not enough—data center operators are also looking for cost-effective, compact form factors with low power consumption. 400ZR transceivers are viewed by many as the right solution to satisfy these requirements, but requires expertise in low-power DSPs, integrated photonics, and advanced packaging.

Since 2011, Acacia has leveraged the benefits of integration to evolve from 100Gbps in a 5”x7” form factor to 200Gbps in a CFP2 form factor (see Figure 5). These advances have been enabled by (1) development of low-power DSPs that integrate more functionality while reducing power; (2) the integration of multiple discrete optical components into a single package using Acacia’s mature and highly integrated silicon single-chip photonic integrated circuits (PICs) that minimize both size and power consumption;

and (3) innovations in packaging/integration to achieve high densities. Acacia is currently leveraging these capabilities and investing in co-packaging technology to develop 400Gbps pluggable transceivers in the OSFP, QSFP-DD, and CFP2 form factors.

Acacia was first to introduce digital coherent optics (DCO) modules in the CFP and CFP2 form factors and we anticipate that our expertise in low-power DSPs and integrated photonics positions us well to develop 400ZR modules that can scale to meet the required volumes.

Acacia is a trusted partner for many carriers, data center operators, and network equipment manufacturers (NEMs) with products already on the market. This positions us well to also become a trusted partner for 400ZR solutions.

To find out more about how Acacia solutions can support your optical networking needs, visit our website or [contact us](#).

Coherent Optics Basics

Intensity modulation/direct detection: Optical signal intensity is modulated. Receiver detects whether light is on or off.

- Non-Return to Zero (NRZ) intensity modulation: Signal intensity varied by switching the laser output between either fully on or fully off. As a result, the transmitter only sends one bit per symbol. The problem is that the laser switching electronics are limited by semiconductor physics. Once the semiconductor electronics are maxed out, the bit rate of NRZ hits a hard limit.
- Four-level pulse-amplitude modulation (PAM4): Varies optical intensity among four, rather than two, levels, sending two bits per symbol. This requires multiple lasers to reach 400Gbps, however. Dividing peak power of the signal into four levels reduces OSNR, which shortens effective transmission distance.

Coherent modulation/coherent detection: Optical signal is modulated signal in both phase and amplitude. The number of bits per symbol using a single laser is increased, reducing the need for multiple laser sources. Optical signal can be modulated with two orthogonally polarized beams, known as dual-polarization (DP), further doubling the amount of data transmitted. Receiver with DSP recovers data from preserved phase and amplitude information.

- Quadrature phase-shift keying (QPSK): Optical signal split, with one part phase shifted relative to the other, and amplitude held constant. Recombining them generates a complex optical signal that encodes two bits of data for every state change of the laser. DP-QPSK doubles the amount of data transmitted.
- Quadrature-amplitude modulation (QAM): Modulates optical signals simultaneously in phase and amplitude. Bandwidth increase depends on levels of amplitude modulation. 16QAM uses four amplitude levels to encode four bits of data per symbol. DP-16QAM doubles the amount of data transmitted.

References

¹ Cisco (November, 2018), "Cisco Visual Networking Index: Forecast and Trends, 2017-2022," Retrieved from: <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.pdf>.

About Acacia

Acacia's innovative silicon-based high-speed optical interconnect products accelerate network scalability through advancements in performance, capacity, and cost. Our silicon photonic PICs, DSP ASICs, and coherent modules inside a variety of network equipment products empower cloud and service providers to meet the fast-growing consumer demand for data.



3 Mill and Main Place, Suite 400
Maynard, MA 01754, USA
info@acacia-inc.com
978.938.4896
acacia-inc.com