Consortium For On-Board Optics
Trends in Optical Networking Communications

cobo

Contents

Figures

CURRENT TRENDS IN OPTICAL NETWORKING

Hyperscale datacenters are driving optical technology development and are an increasingly important part of the overall optical communications market. According to Synergy Research Group, (https://www.srgresearch.com/articles/hyperscale-operator-spending-data-centers-11-2019-despite-only-modest-capex-growth) there are 420 hyperscale datacenters in the world with another 132 developed by the end of 2019. The United States dominates the market owning 40% of the world's hyperscale datacenters and investments by companies like Microsoft, Google and Amazon are significant. Hyperscale capital expenditure (CapEx) growth has been steady for several years and in 2018 topped $26 billion - up 53% year-over-year.

The key trend driving hyperscale growth is the transition of services and applications out of personal and corporate environments and into the cloud. Development and deployment of hyperscale networks are not based on spanning tree architecture typically used in Enterprise networks; instead, mesh architectures are utilized to provide greater reliability. The mesh network with its leaf and spine switches drive an increase in volume of optical interconnects. As the speed of mesh networks increase, there is growing pressure to utilize optical interconnects all the way to the server.

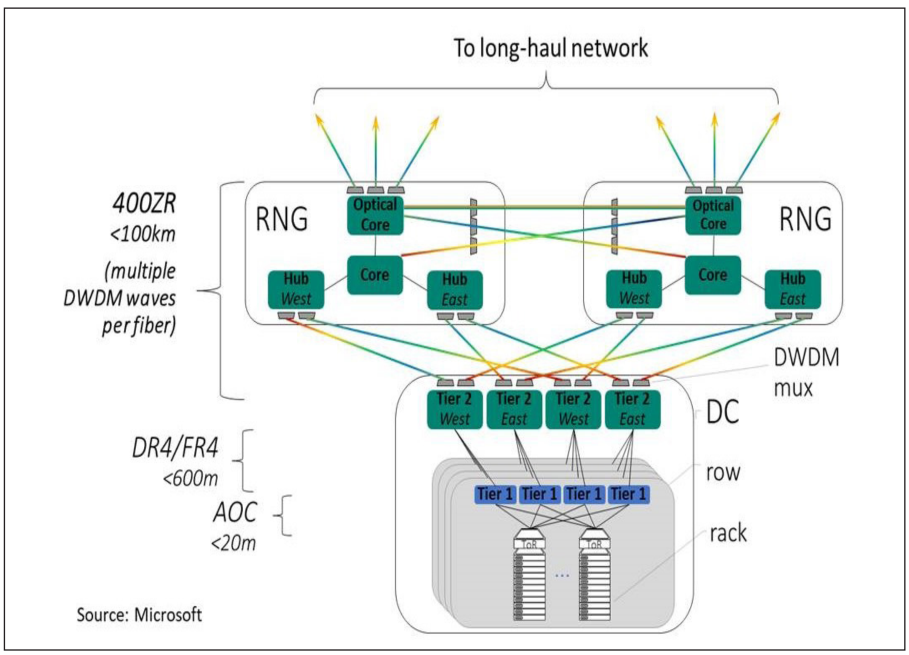

FIGURE 1: DATA CENTER NETWORK INTERCONNECT EXAMPLE

Interest in On-board Optics and Co-packaged Optics (the integration of optics with switch chips) will increase as the number of I/Os expand and as face plate density becomes more important. As interconnect applications move beyond 100G, coherent optics will play a larger role in data centers and these trends will lead to optical technologies in all form factors and speeds.

For example, Active Optical Cable (AOC) deployment is rapidly growing and have recently become popular in hyperscale, enterprise and storage systems as a high-speed, plug & play solutions displace Direct Attach Copper (DAC) cables. AOCs achieve high data rates over long reaches and use a fraction of the power when utilized to connect servers to switches at 25G rates, and between servers and storage at 100G. The AOC market breakdown is shown in Figure 2.

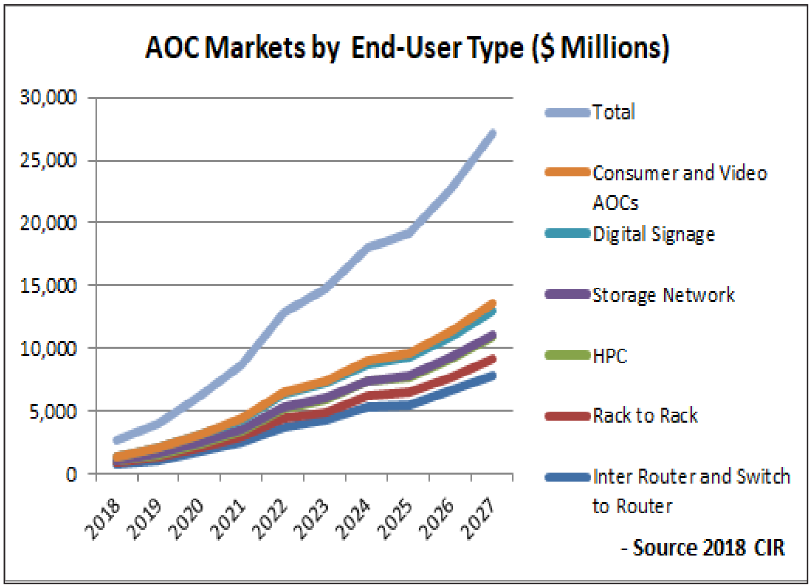

FIGURE 2: MARKET FORECAST FOR AOCS

The Transition to 400G

The transition to the 400G Ethernet ecosystem began in 2019 as switches are able to support 400G ports. Initially 100G Ethernet switches were based on 3.2 Tbps switch silicon and 32 ports of 100G using QSFP28. The transition to 400G Ethernet generation switches is a quadrupling of speed capability as 12.8 Tbps switches are equipped with 32 ports. The big difference at 400G is the availability of new optical form factors and it is anticipated large hyperscale cloud providers will lead the ramp-up in the next few years.

There are three faceplate pluggable 400G transceiver form factors in the market. CFP8 was the first-generation form factor used in early core routers and DWDM transport client interfaces, however, was not targeted for use in the high-volume intra-datacenter applications. For datacenter switches, there are two faceplate pluggable form factors: QSFP-DD and OSFP.

Faceplate pluggable transceivers have increased in bandwidth density per unit of volume at 10/25/50/100G and as bandwidth increases, the pitch remains basically the same. The market is demanding faceplate pluggables to grow bandwidth beyond 400G and supporting on-board optics and co-packaged optic designs provides a solution for continuing the growth curve.

Pluggability Requirements

Embedded optics have been deployed for generations and the concept of on-board or embedded optics is not new in the networking industry. Embedded optics are mounted directly on the host PCB and in supercomputers, are inserted into a socket that is mounted on the host PCB. The previous generations of embedded optics were based on various MSAs or proprietary solutions, and many were targeted for very specific applications.

The Consortium for On-Board Optics (COBO) members developed an embedded optical module form factor specification to provide multi-generation capability, broader technology capability, and enhanced thermal and signal integrity properties. The COBO specification for 8-lane and 16-lane on-board optical modules is an open specification offering module and system designers to use one of the two footprint widths to support the three module lengths.

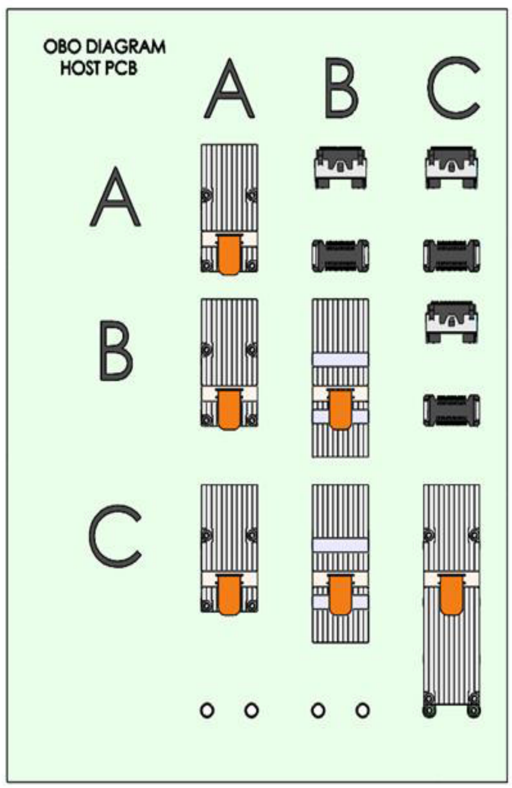

FIGURE 3: HOST-OBO CLASS COMPATIBILITY CHART

The industry view today is that for 400G Ethernet, pluggable optics will still be preferred for datacenters operating with a "pay as you grow" strategy or where the facilities maintenance model is harder to modify. The handling of pluggable optics does come at a cost for datacenters. While optics have become the greatest portion of CapEx in building datacenter networks, they also have a significant operational expense (OpEx). For example, switches utilizing faceplate pluggable optics arrive at the datacenter unpopulated, and technicians and operators are required to populate and verify the optics. Datacenters using on-board optics can have these functions performed by their system integrator or switch vendor.

Initial 400G switch electrical interfaces to the optical module are the same as 100G solutions and are based upon 25 Gbaud PAM4 signaling, except utilizing two bits per symbol instead of one. The use of PAM4 signaling reduces the signal-to-noise (SNR) margin available; therefore, forward error correction (FEC) techniques are used to improve the error performance. FEC doesn't solve all SNR issues, and there is still a limitation on the reach and performance of the PCB channel between the switch silicon and the optical module. With faceplate pluggable optics, those requirements create mechanical constraints increasing thermal density and impact air-flow.

Many system vendors have developed solutions to address these concerns with faceplate pluggable optics at 400G and believe they can transition modules to support 800G of bandwidth. The COBO form factor provides a tool in the system development kit for cases where faceplate pluggable optics are not a viable solution and paves the way to 1.6T module bandwidth (or 1.6T Ethernet) that require on-board optics.

COBO Form Factor Breadth of Capabilities

One primary reason datacenter operators deploy faceplate pluggable optics is the flexibility of a single form factor for direct-attach copper cables, active optical cables, multimode optics and single-mode optics. With 100G Ethernet, a single form factor is capable of reaching from less than a meter to over 80 km. The QSFP-DD and OSFP pluggable form factors have been architected to support this wide diversity in reach including the support of the OIF's 80km 400ZR project.

The COBO form factor also supports the OIF's 400ZR application but approached the challenge differently. Instead of a "one size fits all" approach, COBO supports three OBO sizes so that applications requiring only short reach optics can use the smaller COBO form factors greatly increasing density while applications requiring the larger 400ZR optics can target the larger COBO form factor. The COBO form factor achieves this by using the same sockets and pin-out assignments but with the high speed and low speed connectors placed at three different separations as specified in the COBO specification.

Trends in the Data Center Interconnect Market – Coherent versus Direct Detection

Coherent interfaces are rising to capture the Data Center Interconnect market especially in the 80 kilometers range, and at 40 kilometers and below while Direct Detection solutions remain the lower cost and lower power solutions. Coherent technology currently is approximately 2x higher in power and costs than Direct Detection and standards organizations will focus on Direct Detection when possible for the shorter reaches.

The OIF developed an implementation agreement for 400GZR and the IEEE task force is targeting shorter reaches.

The P802.3cn is handling ER8 and the CT task force is addressing ZR interfaces using coherent technology for 80 kilometers applications which leads future innovation to support 1.6Tbps coherent interfaces in 2024. 100G + 200G links at 80-100 kilometers are using Coherent technology requiring a transponder box, most of which are using pluggable coherent CFP2 Analog Coherent Optic and DCO Transceivers, as Analog Coherent Optics are a linear interface between transceiver and the host.

COBO expects transceiver evolution will be driven by optical packaging and DSP power dissipation to enable higher density optical designs to transition to 400G CFP-2 DCO and QSFP-DD DCO Modules plugged directly into switches and routers without the need for an on-board DSP.

Beyond 400GbE

An interface between Top of Rack and Tier 1 and Tier 2 will not fundamentally change the architecture in data centers and the questions remains when and if the front panel pluggable approach fails to be feasible or cost effective.

On-board mounted optical assemblies and co-packaged optics could fill this need and reshape the Optical Business Model at the component and system level if either provides a more cost-effective solution compared to pluggable optics.

The higher capacity solutions relay on a 100G electrical interface to the host which are being developed by the IEEE's 100G PAM4 electrical 802.3ck task force to support the 100G/λ optical standards being defined by the IEEE 802.3cu effort.

Moving Forward

Debates continue on industry profitability, the end of pluggables, and the beginning of commercialized, high volume integrated optic designs. High-volume mature products are attractive for contract manufacturers and new high-performance products offer an opportunity for companies with strong R&D teams and expertise in chip design and manufacturing to benefit at 800G and beyond.

COBO is well positioned to develop the next generation of dense on-board or co-packaged optic standards to support future generations of dense high-capacity optical solutions.

Limitations

THIS DOCUMENT IS PROVIDED "AS IS." The publisher and contributors expressly disclaim any warranties (express, implied, or otherwise), including but not limited to implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title. The entire risk of using or implementing this document is assumed by the user. IN NO EVENT IS THE PUBLISHER OR ANY CONTRIBUTOR LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING LOST PROFITS) BASED ON ANY THEORY OF ACTION (INCLUDING NEGLIGENCE), EVEN IF THE OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. This document may be subject to third party rights.