Data Center Cable Infrastructure

New methods for testing and installing AOC and DAC

By: Guylain Barlow, Senior Product Manager, Viavi Solutions

The data center is where a massive amount of information is being exchanged and where a key requirement is the capacity to scale. Beyond the need for power, cooling, storage and switching inside the data center is the necessity for practical and efficient cabling. Data centers can roughly be divided between hyperscale, multi-tenant, and private. The use of cables as described in this paper, namely Active Optical Cables (AOC) and Direct Attach Copper (DAC), applies to all three categories, although this topic is especially applicable to hyperscale data centers. This application note covers practical operational considerations such as the validation of AOC and DAC to save time and reduce costs in data centers.

Data Center Architectures

Figure 1 provides an example of a data center and its interconnectivity to the external world. Within the data center, there are a few potential architectures:

Top of Rack (TOR) architecture is where the cabling between switch and server stays within a rack. This has the benefit of reducing the overall amount of cabling with the downside of reduced efficiency in the usage of Ethernet switch ports which are limited to within a rack.

End of Row/Middle of Row (EOR/MOR) configuration is where switch ports are grouped together leading to longer cables. There are two EOR/MOR examples where in one case, cables connect directly between servers and switch ports. In the second case, physical connectivity goes through a patch panel which provides the benefit of greater connection flexibility with the disadvantage of a higher number of cables.

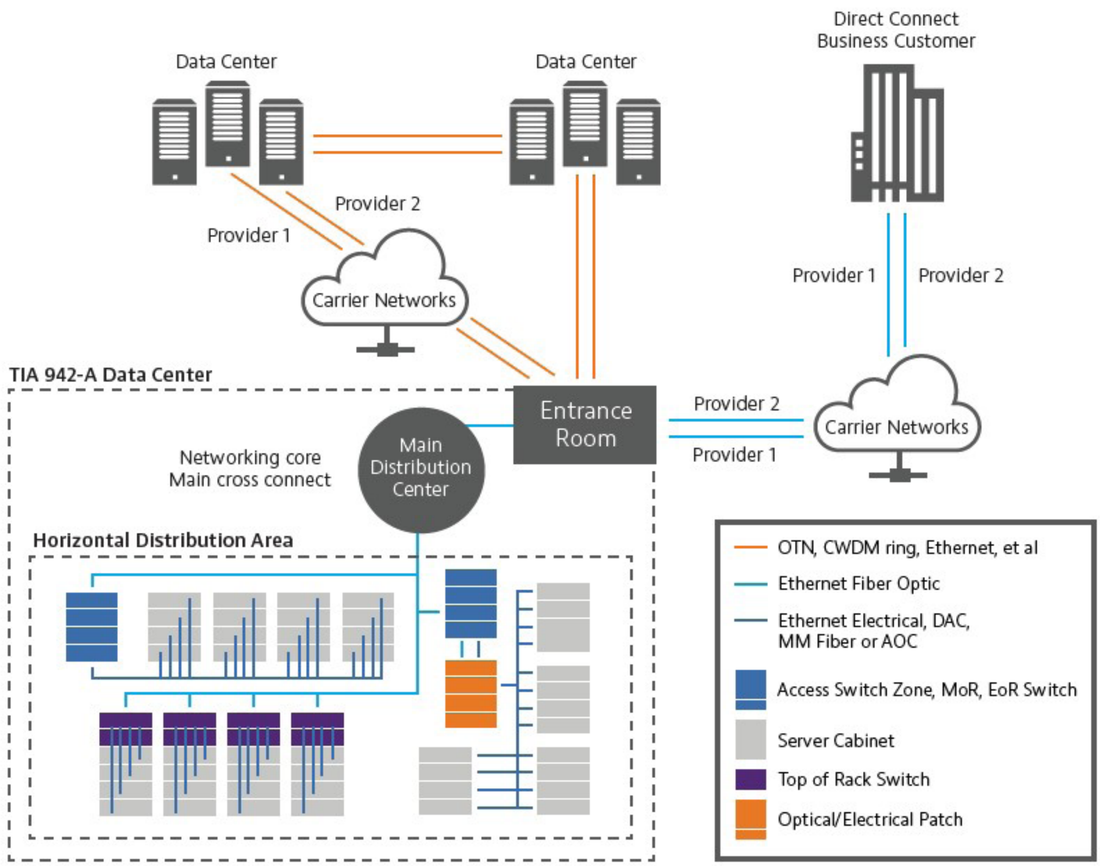

Figure 1. Data Center Architecture

Active Optical Cables

Active Optical Cables as shown in figure 2 are used in limited range interconnect applications in data centers. For high-speed links at 40GE and 100GE, this means using multiple lanes of data over ribbon cables. At 10GE or 25GE, a single lane or fiber per direction is sufficient. An AOC is often based on multi-mode fibers while some (like PSM4) is on single-mode fiber. A key attribute is that AOC employs the same connectors as pluggable optics and performs electrical-to-optical conversions at each cable end. In practice, this means QSFP terminations for 40GE and 100GE and SFP terminations for 10GE and 25GE. The AOC is therefore active and includes transceivers, control chips and modules, in addition to the fiber optics cable. AOC cables are of fixed length, starting at just a few meters with possibilities up to 100 meters or more. Technically, an AOC does not have to comply with an Ethernet interface type although many advertise a certain type in their coded information; table 1 provides a list of potential Ethernet interface types. From the table, RS-FEC stands for Reed-Solomon Forward Error Correction; it is a digital mechanism designed to extend transmission distance by adding redundancy to a signal which enables code word self-correction at the far-end. The RS-FEC algorithm, when specified for use with a cable, runs on the Ethernet switches and servers found at each end of the physical connection.

Figure 2. AOC Cable

Table 1 Ethernet Interface Types

Ethernet Rate	Interface Type	Maximum Reach	Medium	No. of fibers/ wavelengths	Wavelength Range	RS-FEC	Pluggable Type
100GE	CWDM4 MSA	2 km	SMF	4 λ / dir	1310 nm	Yes	QSFP28
	CLR4 Alliance	2 km	SMF	4 λ / dir	1310 nm	Optional	QSFP28
	PSM4 MSA	500 m	SMF	4 fibers / dir	1310 nm	Yes	QSFP28
	SWDM4 Alliance	75 m 100 m 150 m	OM3 MMF OM4 MMF OM5 MMF	4 λ / dir	850 nm	Yes	QSFP28
	100GBASE-SR4	70 m 100 m	OM3 MMF OM4 MMF	4 fibers / dir	850 nm	Yes	QSFP28
	100GBASE CR4	5 m	Twinaxial	4 cables / dir	N/A	Yes	QSFP28
40GE	SWDM4 Alliance	240 m 350 m 440 m	OM3 MMF OM4 MMF OM5 MMF	4 λ / dir	850 nm	No	QSFP+
	40GBASE-SR4	100 m 150 m	OM3 MMF OM4 MMF	4 fibers / dir	850 nm	No	QSFP+
	40GBASE CR4	7 m	Twinaxial	4 cables / dir	N/A	No	QSFP+
25GE	25GBASE-SR	70 m 100 m	OM3 MMF OM4 MMF	1 fiber / dir	850 nm	Yes	SFP28
	25GBASE CR	5 m	Twinaxial	1 cable / dir	N/A	Yes	SFP28
	25GBASE CR-S	3 m	Twinaxial	1 cable / dir	N/A	No	SFP28
10GE	10GBASE-SR	33 m 400m	62.5 μm MMF 50 μm MMF	1 fiber / dir	850 nm	No	SFP+
	10GBASE- CR	15 m	Twinaxial	1 cable / dir	N/A	No	SFP+

Direct Attach Copper Cables

Figure 3. DAC Cable

Direct Attach Copper (DAC) cables as shown in Figure 3 are an alternative when the cable itself is made of copper instead of an optical fiber. A DAC may be passive to provide a direct electrical connection or active when signal processing circuitry is integrated in the DAC built-in connectors. Just as with an AOC, a DAC will be terminated by SFP or QSFP depending on the line rate. As a comparison, AOC cables support longer transmission distances, use less power, and are more lightweight than DAC cables. However, they cost more and optical fibers can be more easily damaged than copper cables. When comparing AOC cables to traditional fiber optic cables connected to pluggable optics, AOCs provide a simplicity of installation without the need to consider interconnection loss, and eliminating of the need to clean and inspect fiber end-faces before making a connection. However, AOC cables cannot be used in EOR/MOR configurations that use patch panels as described earlier.

Operational Challenges

Because AOC and DAC cables do not provide test access to the actual fiber or copper cabling, traditional media test and certification tools cannot be used to certify the cable. Instead, a test tool that can accept dual SFP/QSFP transceivers and generate and analyze traffic must be used. Testing AOC and DAC is a critical step to ensure any issues with network performance are not due to the AOC/DAC cable or its installation. Consider that it is costlier to troubleshoot a faulty cable once installed rather than testing it upfront. For one thing, it is necessary to trace and locate the far end. AOC/DAC cable failure causes include simple manufacturing defects with wrong or reversed polarities, or mislabeling or damage during shipment. For AOC, they may get excessively bent causing high loss or the fibers may get crushed. In the case of DAC, there can be EMI degradation resulting in excessive bit errors. Prior to installation, a technician can choose to test all AOC/DAC cables using a test device equipped with dual-port SFP/QSFP or simply sample test a subset of the cables received from a batch. Troubleshooting cables that are already installed requires two devices simply because of the distance between the two cable connector endings.

Bit Error Rate Testing

The simplest and most effective way to test cables is to run a test pattern where the results can be compared to a Bit Error Rate (BER) threshold. AOC and DAC cables usually have a BER rating on their datasheets, especially when they are meant to be used with devices implementing the RS-FEC algorithm. This BER rating depends on the type of cable, line rate, and type of Ethernet interface. In the case of a cable meant for use with RS-FEC encoded traffic, which is typical at 100GE and 25GE, there may even be both a pre-FEC rating, which is before error correction, and post-FEC rating after error correction. In such a case, it is recommended to perform a cable test using a pre-FEC BER threshold close to the cable BER rating and ensure the measured BER is smaller than the threshold for a successful test. For 40GE and 10GE cables where RS-FEC is not used, the expected BER threshold needs be quite a bit smaller as there is no error correction on those circuits. In such cases, if there is no BER rating for the AOC or DAC, the recommended threshold BER is 10^-12. Test times of one minute per cable are more than sufficient to obtain meaningful BER results with line rates of 10Gbps or higher. Best practice procedures for cable tests will result in the generation of test reports including a cable identifier, such as the serial number, which can be read from an AOC or DAC cable. In summary, ensuring AOC or DAC cables are tested against their target BER thresholds is a meaningful method to ensure cables will be functional when connected to switches and servers in the data center.

Viavi T-BERD/MTS 5800–100G

To test AOC and DAC cables, Viavi has introduced an integrated script to automate cable assembly testing. Viavi's T-BERD/MTS 5800-100G, shown in Figure 4, provides all-in-one testing at line rates up to 112Gbps. The 5800-100G supports all Ethernet rates with dual-port capabilities including 10/100/1000BASE-T, optical GE, 10GE, 25GE, 40GE, and 100GE. Technicians can test numerous applications including AOC/DAC cables in addition to metro, backbone, and data center interconnectivity. Although a small size, the 5800-100G can test from DS1/E1 to OTU4 including CPRI, Fiber Channel, PDH, SONET/SDH, OTN, and Ethernet. In addition, the 5800 platform expands to support OTDR modules, fiber inspection with auto-focus, and advanced timing and synchronization capabilities. These tools are all supported by StrataSync which provides cloud-based asset management, configuration, and reporting for easy management and sharing. The 5800-100G is a complete test and measurement solution to support the needs of the data center user in the data center , including AOC and DAC cable testing.

Figure 4. T-BERD/MTS 5800-100G

Contact Us **+1 844 GO VIAVI**
(+1 844 468 4284)

To reach the Viavi office nearest you,
visit viavisolutions.com/contacts.

© 2017 Viavi Solutions Inc.
Product specifications and descriptions in this document are subject to change without notice.

dc-cable-infrastructure-an-fit-nse-ae
30186165 000 0717