

Application Note

Maintaining Robust DAC, AOC, AEC and Transceiver Connectivity in Hyperscale Data Centers

A case for automated, multi-layer testing

Exponential growth in data creation, storage and processing is driving global demand for hyperscale data centers. The infrastructure of these data centers, most of which is designed and managed by cloud service providers (CSPs) such as Amazon, Microsoft, and Google, relies on robust physical connectivity to ensure that deployed storage and compute resources are fully utilized. This application note makes a case for using an automated, multi-layer test solution to certify and troubleshoot these critical physical connections.

The three most common physical links found within hyperscale data centers are: Direct Attach Copper (DACs) cables, Active Optical Cables (AOCs) and now Active Electrical Cables (AECs). All three of these connectivity elements can be viewed as highways over which data travels. Not only is it important that the underlying electrical and optical connectivity of these components function properly, it's also critical that data transmission (through various communication protocols) operates within the established network/matrix specifications.

These connectivity elements vary primarily by data rate and physical link length (maximum reach). DACs are used to connect devices separated by short distances and are most often leveraged to connect two elements within the same rack. More recent AEC (Active Electrical Cable) assemblies extend distance coverage to copper cables by adding active elements in each cable. AOCs have extended reach that enables the connection of two devices within the same row or even to a neighboring row. Fiber optic cable assemblies connecting to pluggable optics or including long-haul spliced connections can connect two devices at great distance from one another (typically up to 80km without amplification depending on the transceivers to which they are connected). No matter what the data rate or distance of span, all of these links must be tested to ensure functional operation.



Copper/Electrical Cables: Passive or Active

A Direct Attach Copper (DAC) cable is typically passive and has pluggable modules fixed at either end of a twinax copper cable [small form factor pluggables (SFPs) or quad small form factor pluggables (QSFP) depending on the line rate]. Both ends have specific connectors, and the cable length is fixed. Electromagnetic (EMI) shielding around the copper cable increases, or thickens, as the speeds increase to keep communication reliable. For longer reaches, copper/electrical cables called AEC (Active Electrical Cable) are active and provide signal processing circuitry integrated in the built-in connectors, this way signals can be reshaped and potentially amplified.

AECs are relatively new and aim to match the transmission rates of AOC cables (at shorter distances), but more economically. AECs are narrow gauge copper cables with industry standard connectors at each end. Its active cables are mainly used for the connection of top of rack servers, distributed chassis, and up to 500 cables per rack. AEC active cable supports transmission rates of 100G, 200G, 400G, optics types include OSFP56, OSFP, OSFP-DD.



Figure 1. DAC Cable



Figure 2. AEC Cable

Active Optical Cables (AOC)

Active Optical Cables (AOC) as shown in Figure 3 and Figure 4 are used in point-to-point and break-out interconnect applications in data centers, most often within the same rack or row. AOC cables were designed to fix capacity and reach problems in the world's data centers. Traditional copper cabling used in data centers was heavy, making moves, adds, and changes very difficult. In addition, EMI given off by data center equipment hampered transmission performance on the copper cables, often significantly. AOCs are naturally immune to EMI, and are also lighter, thinner, and have a longer range than the DAC and AEC. For example, an AOC cable can transmit at 400Gbps for 100 m compared to only 2 m for DAC and 7m for AEC. The AOC is actually an active assembly which explains its greater range and higher cost, and is comprised of transceivers, control chips, modules fused to either end, and the fiber optics cable. When selecting cables, the length, data rate, and transceiver type must be chosen with some certainty because they cannot be changed once the AOC is manufactured. Figure 4 shows an example of breakout cable (DACs and AECs also have breakout options) such as 4x10GE, 4x25GE, or potentially 4x100GE.



Figure 3. AOC Cable




Figure 4. AOC Breakout Cable

Table 1. Ethernet Interface Types

Ethernet Rate	Interface Type	Maximum Reach	Medium	No. of fibers / wavelengths	Wave length Range	RS-FEC*	Pluggable Type
400GE	400GBASE-DR4	500 m	SM	4 fibers/dir	850 nm	Yes RS (544, 514)	QSFP56-DD or OSFP
	400GBASE-SR8	70 m	OM3 MMF	8 fibers/dir	850 nm	Yes RS (544, 514)	QSFP56-DD
		100 m	OM4 MMF	8 fibers/dir	850 nm	Yes RS (544, 514)	QSFP56-DD
	400GBASE-SR4.2	70m	OM3 MMF	8 fibers 2 λ /dir	850 nm	Yes RS (544,514)	QSFP56-DD
		100 m	OM4 MMF	8 fibers 2 λ /dir	850 nm	Yes RS (544,514)	QSFP56-DD
		150 m	OM5 MMF	8 fibers 2 λ /dir	850 nm	Yes RS (544,514)	QSFP56-DD
400GBASE-CR8	3 m	Twinaxial	8 cables/dir	N/A	Yes RS (544, 514)	QSFP56-DD	
200GE	200GBASE-SR4	70 m	OM3 MMF	4 fibers/dir	850 nm	Yes RS (544,514)	QSFP56
		100m	OM4 MMF	4 fibers/dir	850 nm	Yes RS (544,514)	QSFP56
	200GBASE-CR4	3m	Coaxial	4 cables/dir	N/A	Yes RS (544,514)	QSFP56
100GE	100GBASE-DR	500 m	SMF	1 fiber/dir	1310 nm	Yes RS (544,514)	QSFP28 or QSFP56
	PSM4 MSA	500 m	SMF	4 fibers/dir	1310 nm	Yes RS (528,514)	QSFP28
	100GBASE-SR4	70 m 100 m	OM3 MMF OM4 MMF	4 fibers/dir	850 nm	Yes RS (528,514)	QSFP28
	100GBASE-CR4	5 m	Twinaxial	4 cables/dir	N/A	Yes RS (528,514)	QSFP28
50GE	50GBASE-SR	70 m	OM3 MMF	1 fiber/dir	850 nm	Yes RS (544,514)	QSFP28 or SFP56
	50GBASE-SR	100 m	OM4 MMF	1 fiber/dir	805 nm	Yes RS (544,514)	QSFP28 or SFP56
	50GBASE-CR	3 m	Twinaxial	1 cable/dir	N/A	Yes RS (544,514)	QSFP28 or SFP56
40GE	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 RS (528,514)	SFP28
	25GBASE-CR	5 m	Twinaxial	1 cable/dir	N/A	Yes RS (528,514)	SFP28
	25GBASE-CR-S	3 m	Twinaxial	1 cable/dir	N/A	No	SFP28
10GE	10GBASE-SR	33 m 400m	62.5 μ m MMF 50 μ m MMF	1 fiber/dir	850 nm	No	SFP+
	10GBASE-CR	15 m	Twinaxial	1 cable/dir	N/A	No	SFP+

*RS-FEC stands for Reed-Solomon Forward Error Correction, which 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.

The Challenge of Testing DAC and AOC Cables

Because DAC, AEC and AOC do not provide test access to the actual fiber or copper cabling, traditional media test and certification tools cannot be used to certify or troubleshoot the cable. Instead, a test tool that can accept dual transceivers and generate and analyze traffic must be used. Testing these high speed cables is a critical step to ensure any issues with network performance are not due to the DAC/AEC/AOC 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. DAC/AEC/AOC 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 or AEC, there can be EMI degradation resulting in excessive bit errors. Add the increased number of cables to be tested in the hyperscale data center and it's easy to understand the need to automate the testing process.

Edge Deployments and Disaggregation – Balancing Time/Cost Efficiencies of Testing at Install

Hyperscale data center construction and commissioning in the era of optical network disaggregation means that contractors responsible are now also responsible for testing and certification of performance, operability, stress and reliability of every fiber link within these multi-component white box networks. Combine this with the exponential growth of hyperscalers and their need to be ever closer to the end user and the result is more edge deployments (network virtualization). This is forcing hyperscalers to increase speed, security and efficiency at the same time as they minimize latency. The coinciding need to turn up edge deployments quickly can add fuel to the decision not to test all cables prior to installation, opting instead to wait and address any connectivity issues during troubleshooting.

Similarly, the need to minimize downtime during troubleshooting often leads to the decision to cut or disconnect a cable and lay a new cable rather than troubleshoot or remove the existing cable. It is common for untested cables pulled out of cabinets in haste to be returned to the manufacturer only to have the manufacturer say there is nothing wrong with the cable, or to collect large volumes of “non-working” cables with no ability to diagnose them. Not only is this costly — cables can vary from tens of dollars to thousands depending on the line rate — leaving dead cables also creates unwanted bulk in the cabinet. This bulk is cumbersome and can lead to mislabeling, confusion and the increased probability of unplugging a live cable. Because old cables cannot be used in an upgrade due to their rate-specificity, leaving cut and dead cables in a cabinet creates unsustainable volume and weight that can compromise the function of the rack and the structure.

The Value of Bit Error Rate Testing

Accurately illustrating the time-cost benefit of testing and verifying every cable at install is difficult due to the varying costs of the cables among other factors. However, it's not difficult to theorize the consequences of how not testing enough cables at install could make future troubleshooting efforts as well as upgrades to the network more time consuming and costly.

The simplest and most cost-effective way to test cables is to run a test pattern where the results can be compared to a Bit Error Rate (BER) threshold. DAC, AEC and AOC cables including breakouts 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 (see Table 1). In the case of a cable meant for use with RS-FEC encoded traffic, which is typical at 400GE, 200GE, 100GE, 50GE and 25GE, there may even be both a pre-FEC rating (before error correction) and a 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 to 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 DAC or AOC, 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 a DAC or AOC cable. Testing DAC or AOC cables against their target BER thresholds is a meaningful method to ensure more cables are functional when connected.

All-in-One Automated Testing and Troubleshooting Solutions for the Hyperscaler

OneAdvisor 800

DAC, AEC, and AOC Cable Test coverage with OneAdvisor 800 is quite complete. This includes the testing of cables such as 4x100G QSFP28 to QSFP-DD, 4x100G QSFP56 to QSFP-DD, and 200G-200G cables with a single instrument. Further the TM400GB-QQ module tests 400G-400G cables with a single module. The VIAVI OneAdvisor 800 with 400G module supports all Ethernet rates from 10/100/1000Base-T up to 400GE. It provides full simultaneous and independent dual-port testing up to 400GE.

T-BERD/MTS-5800 Automated Cable Test Suite

Illustrated in Figure 5, the VIAVI T-BERD/MTS-5800 platform is a suite of instruments that support the multi-rate testing and troubleshooting of DACs, AOCs, and transceivers in hyperscale data center networks. This platform helps data center ops technicians to reduce cost and increase network uptime through fast, automated testing. VIAVI provides the cable test integrated script to automate DAC/AOC/transceiver testing. Cable test supports the testing of either active/passive DAC or AOC in addition to transceivers. Integrated into the VIAVI cable test is the unique troubleshooter function which provides an algorithm that looks for optimal cable parameter settings when a typical test yields worse than expected results. Such a breakout cable can be tested with a single unit from each termination point and the job manager tool wraps up the results in a single test report file. VIAVI T-BERD/MTS-5800-100G, 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, 50GE and 100GE.

Prior to installation, a technician can choose to test all DAC/AOC/AEC cables using a single test device equipped with dual-port SFP/ QSFP, or sample test a subset of the cables received from a batch.

During Troubleshooting, cables that are already installed typically require using two test devices because of the distance between the two cable connector endings. In a Top of Row (TOR) configuration, the DAC or AOC can be tested using a single dual-port test device. Technicians can test numerous applications including DAC/AOC cables in addition to metro, backbone and data center interconnectivity.

Extended T-BERD®/MTS-5800 Product Family

Comprehensive Coverage

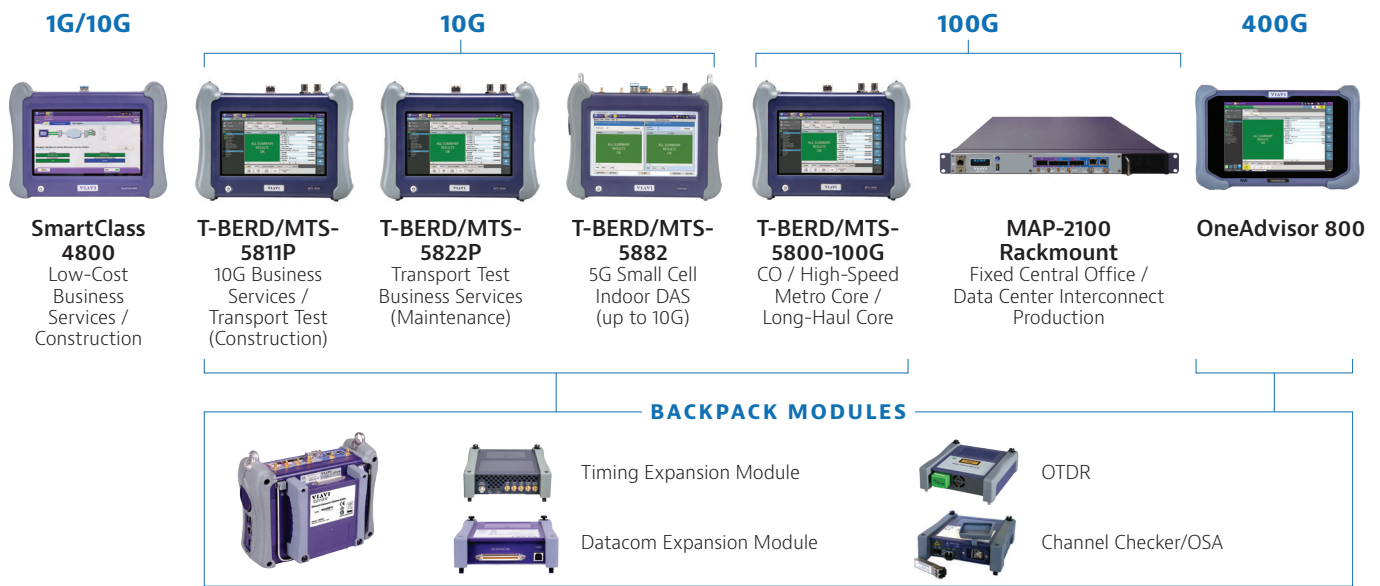


Figure 5. T-BERD/MTS 5800 Automated Cable Test Suite

Benefits of Using the VIAVI T-BERD/MTS-5800 Automated Cable Test Suite:

- Increase speed of construction, commissioning and de-commissioning
- Increase speed of data center monetization
- Reduce time and costs of install, verification, turn-up and troubleshooting
- Improve network performance and data center availability/uptime
- Support SLA delivery