



Optimizing New Short-Reach Singlemode Deployments with Verified Ultra-Low Loss (ULL) Components

Author: Gary Bernstein, RCDD, CDCD



Reading Time: 18 minutes

Growth of Singlemode in the Data Center

As global digitization continues to expand, data centers that are at the core of every organization's critical operations are demanding ever-increasing bandwidth. To meet the demand, advancements in signaling, transceiver, and optical technology and ensuing IEEE fiber optic Ethernet standards have continued to keep up. Multimode and singlemode fiber applications up to 400 Gigabit per second (Gb/s) are now available, and standards for 800 Gb/s Ethernet are already well under development, with future 1.6 and 3.2 Terabit per second (Tb/s) applications not far behind.

Due to its overall greater bandwidth potential and transmission distance, singlemode fiber has rapidly become the media type of choice in hyperscale and cloud data centers for today's higher-speed applications. This adoption has significantly increased purchase volumes and helped drive down the cost of singlemode transceivers. However, for most enterprise data centers where link lengths are shorter, it does not economically make sense to implement expensive high-power lasers used in long-distance 10-kilometers (km), or greater

singlemode deployments. At the same time, cost-effective multimode applications beyond 40 Gb/s are limited to distances of 100 m over OM4 multimode fiber and 70 m over OM3, which can present a challenge for data center designers.

Cost-effectively supporting higher-speed applications to distances beyond 100 m is now feasible with short-reach singlemode applications that support 100, 200, and 400 Gb/s speeds to 500 m and 2000 m distances. However, these new applications also come with significantly reduced insertion loss budgets compared to long-distance singlemode applications and include some return loss considerations. Thankfully, just like they do with multimode applications, verified ultra-low loss (ULL), high-performance cabling systems give data center owners and operators the confidence to reliably support short-reach singlemode applications with margin, while enabling multiple connections for flexibility, manageability, scalability, and speed of deployment.

What is Short-Reach Singlemode?

Long-distance singlemode applications, such as long reach (LR) and extended reach (ER), use high-power edge-emitting laser diodes to deliver more bandwidth over much further distances, with LR applications supporting distances up to 10 kilometers (km) and ER supporting distances up to 40 km. Due to the complex structure of these lasers, long-distance singlemode transceivers cost significantly more

than multimode transceivers that use simpler vertical-cavity surface-emitting lasers (VCSELs).

With cost-sensitive enterprise data centers containing much shorter link lengths, multimode fiber became the media of choice over the past several decades. However, as transmission speeds have increased, the lengths supported by multimode fiber have decreased. For example,

the 10 Gb/s 10GBASE-SR duplex application can be supported up to about 550 m over OM4 multimode fiber. For 40 Gb/s 40GBASE-SR4 parallel optic applications that use multiple fibers with multi-fiber push-on (MPO) connectivity, supported distances drop to just 150 m over OM4 multimode. For 100 Gb/s multimode applications and beyond, the OM4 distance is limited to 100 m.

To support longer distances for high-speed data center links, power-conserving singlemode lasers with relaxed specifications were developed for short-reach data center deployments. With less power and associated energy consumption and heat generation, these transceivers are less expensive than long-distance singlemode lasers. The preliminary options for short-reach singlemode applications were based on non-return-to-zero (NRZ) signaling schemes at 25 Gb/s per lane. A 100 Gb/s coarse wavelength division multiplexing

(CWDM) application was published by the CWDM4 multi-source agreement (MSA) in 2014. The application, referred to as 100GBASE-CWDM4, features four lanes of 25 Gb/s multiplexed on a duplex singlemode fiber. A 100 Gb/s parallel singlemode four-lane (PSM4) application was also defined in 2014 by the 100G PSM4 MSA, featuring four lanes of 25 Gb/s over eight fibers using MPO connectivity.

With the development of the four-level pulse amplitude modulation (PAM4) signaling scheme that offers twice the bit rate of previous NRZ signaling to enable 50 and 100 Gb/s per lane, the IEEE introduced data center reach (DR) and fiber reach (FR) singlemode applications with distances of 500 m and 2000 m, respectively. PAM4 signaling paved the way for cost-effective 100 Gb/s duplex and four-lane 200 and 400 Gb/s short-reach singlemode applications. Table 1 shows the various options for short-reach singlemode applications.

Application	Standard	Distance	Signaling/ Lane Speed	Number of Lanes	Transmission	Number of Fibers
100GBASE-PSM4	PSM4 MSA	500 m	NRZ 25 Gb/s	4	Parallel Optics (MPO)	8
100GBASE-CWDM4	CWDM4 MSA	2 km	NRZ 25 Gb/s	4	CWDM	2
100GBASE-DR	IEEE 802.3ba	500 m	PAM4 100 Gb/s	1	Duplex	2
200GBASE-DR4	IEEE 802.3bs	500 m	PAM4 50 Gb/s	4	Parallel Optics (MPO)	8
200GBASE-FR4	IEEE 802.3ba	2 km	PAM4 50 Gb/s	4	WDM	2
400GBASE-DR4	IEEE 802.3bs	500 m	PAM4 100 Gb/s	4	Parallel Optics (MPO)	8
400GBASE-FR4	IEEE 802.3cu	2 km	PAM4 100 Gb/s	4	WDM	2

Table 1: Short-reach singlemode applications.

The introduction of PAM4, combined with increased singlemode purchase volumes among hyperscale and cloud data centers, has driven down the cost of singlemode fiber systems. On average, a short-reach 100GBASE-DR system now costs 80% less than a CWDM4 system and about 10% less than a PSM4 system.

As shown in Figure 1, the cost of a 100GBASE-DR singlemode system is also now on par with or less than 100 Gb/s multimode applications, including the 100GBASE-SR4 parallel optic multimode application and 100 Gb/s bi-directional (BD) multimode applications.

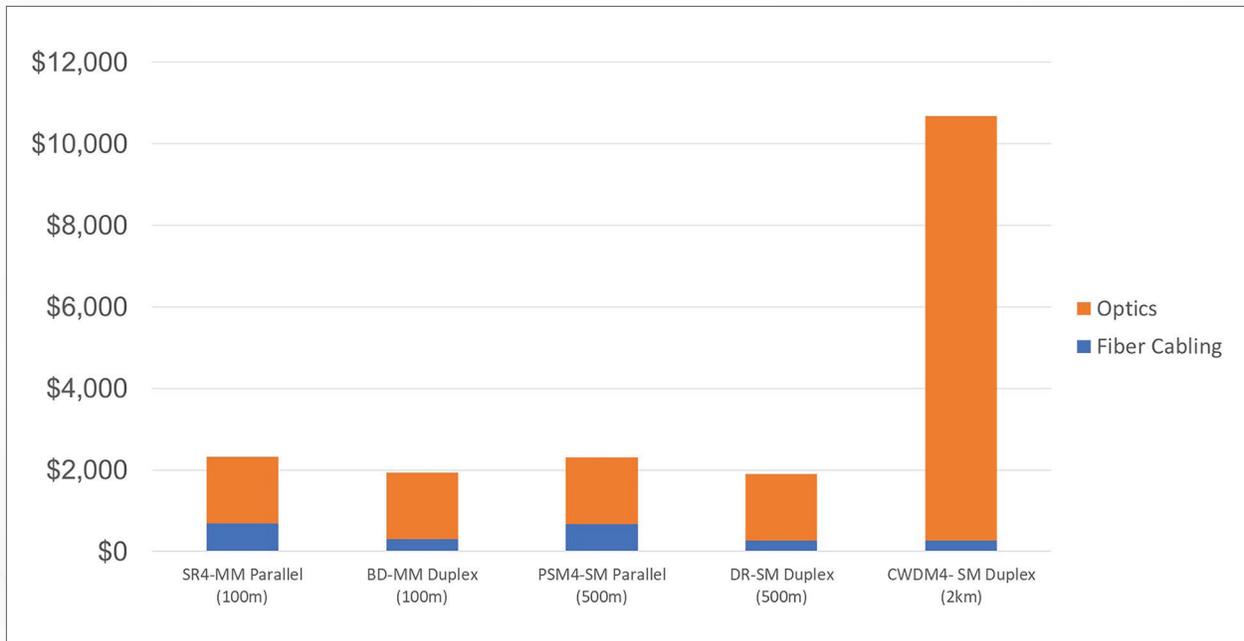


Figure 1: 100G channel pricing, based on average list pricing.

Due to its cost savings, short-reach DR and FR singlemode applications are expected to rapidly gain ground in enterprise data centers as they migrate to 100, 200, and 400 Gb/s speeds in backbone switch-to-switch links. They also support cost-effective breakout configurations where a single 200 Gigabit switch port breaks out to four 50 Gb/s connections or a single 400 Gigabit switch port breaks out to four 100 Gb/s connections, thereby optimizing port utilization and switch density for a reduced cost.

DR and FR applications are also well-positioned for cloud and hyperscale data centers that are migrating to 400 and 800 Gb/s backbone switch-to-switch links and 50, 100, and 200 Gb/s horizontal switch-to-server links. The IEEE 802.3df 400G/800G/1.6T task force has included DR and FR applications as part of the physical layer objectives for 800 Gb/s and 1.6 Tb/s applications, which will further drive the adoption of these applications in cloud and hyperscale data centers.

New Insertion Loss and Reflectance Requirements

Measured in decibels (dB), insertion loss is the amount of power that a signal loses (i.e., attenuates) as it travels along a cable link, which is impacted by the length of the cable and any connection point. Insertion loss can also be impacted by installation variables such as exceeding the bend radius of the cable and fiber connector end face misalignment or contamination, as well as transceiver age. Insertion loss is the primary performance parameter in fiber optic applications – if it is

too high, the signal cannot be properly interpreted by transceivers at the far end. Maximum channel insertion loss is determined by transceiver manufacturers and supporting standards bodies for all fiber applications. As a requirement for Tier 1 and Tier 2 fiber optic testing following installation, insertion loss is measured using a power meter or optical loss test set that compares power injected at one end of the channel to the power received at the other end.

Application	Distance	Max Channel IL
Multimode		
SR/SR4	100 m	1.9 dB
SRBD (BiDi)	100m	1.9 dB
Short-Reach Singlemode*		
DR/DR4	500 m	3 dB
100GBASE-CWDM4	2 km	5 dB
FR4/FR8	2 km	4 dB
Long-Reach Singlemode*		
LR/LR4	10 km	6.3 dB
ER4/ER8	40 km	15 – 18 dB

Table 2: Maximum channel insertion loss requirements.

*Includes 0.2dB fiber loss for 500 m and 0.8 dB for 2 km

While long-distance singlemode LR and ER fiber applications offer ample insertion loss budgets that can reach 18 dB, enterprise data centers have long been challenged by the far more stringent insertion loss budget of just 1.9 dB for 40 Gb/s and faster multimode fiber systems. The lower maximum insertion loss requirements for high-speed multimode systems were a key driver in the development of low-loss

and ULL multimode connectivity. As shown in Table 2, the relaxed specifications for newer short-reach DR and FR singlemode applications also come with significantly reduced insertion loss budgets compared to long-haul singlemode applications. However, at 3 to 5 dB, short-reach singlemode maximum channel insertion loss requirements are still less restrictive than multimode applications.

Another key performance parameter in fiber optic systems is reflectance, which is the amount of light reflected by a connector compared to the amount of light injected. Reflectance is caused by impurities in the fiber and changes between media with different refractive indexes, such as glass to air, and it occurs at mated, open, and contaminated connectors. Like insertion loss, reflectance is also measured in dBs, but it is expressed as a negative number. Note that while IEEE application standards specify reflectance, cabling standards like ANSI/TIA-568-3.D refer to return loss, which is expressed as a positive number. Return loss is simply the inverse of reflectance and is defined as the amount of light injected compared to the amount of light reflected back. A simple rule to remember is that the further reflectance or return loss value is from zero, the better.

Singlemode transceivers are more susceptible to reflectance, especially at wavelengths above 1500 nanometers (nm). That is why angled physical contact (APC) connectors are typically used in singlemode applications and are the de facto standard in service provider FTTX and passive optical networks that send video over higher wavelengths on singlemode fiber. Unlike ultra physical contact (UPC) connectors that feature a rounded fiber end face, APC connectors are polished at an 8-degree angle.

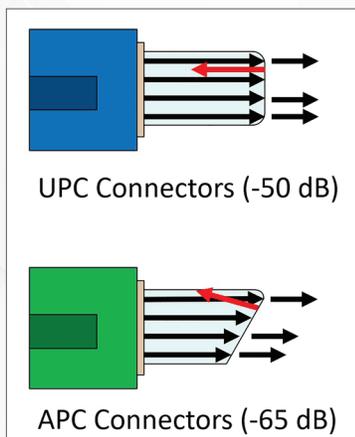


Figure 2: APC connectors with an 8-degree angled end face reflect more light into the cladding, reducing the amount reflected to the transceiver.

The angled end face causes much of the reflected light signal to be deflected into the fiber cladding that surrounds the fiber core, reducing the amount of light reflected to the transceiver. As shown in Figure 2, UPC connectors typically have a reflectance value of around -50 dB, while APC connectors have a reflectance value of about -65 dB. Per industry standards, singlemode UPC connectors are denoted by a blue connector body, while APC connectors are green.

400GBASE-DR4 (IEEE802.3bs)	
# of Discrete Reflections Above -55 dB	Maximum Value for Each Discrete Reflectance
1	-37 dB
2	-42 dB
4	-45 dB
6	-47 dB
8	-48 dB
10	-49 dB
400GBASE-FR4 (IEEE802.3cu)	
# of Discrete Reflections Above -55 dB	Maximum Value for Each Discrete Reflectance
1	-25 dB
2	-31 dB
4	-35 dB
6	-38 dB
8	-40 dB
10	-41 dB

Table 3: Because high-speed short-reach singlemode transceivers are more susceptible to reflectance, IEEE specified reflectance values for DR and FR applications based on the number of mated pairs in the channel.

High-speed short-reach singlemode transceivers are even more susceptible to reflectance. That caused IEEE to specify reflectance values for DR and FR applications based on the number of mated pairs in the channel. As shown in Table 3, if the reflectance of mated connectors in a 400GBASE-DR4 application is -45 dB, only four connectors can be included in the channel. With an improved reflectance of -49 dB, 10 connectors can be included. If the discrete reflectance requirements for DR and FR applications cannot be met, the maximum channel insertion loss will need to be reduced.

The Benefits of ULL Connectivity

While newer short-reach singlemode applications are a cost-effective option for high-speed data center links up to 2000 m, data center designers need to consider the more stringent insertion loss requirements

and carefully review the loss values of fiber components, including jumpers, trunks, modules, and adapters used in common data center configurations shown in Figure 3.

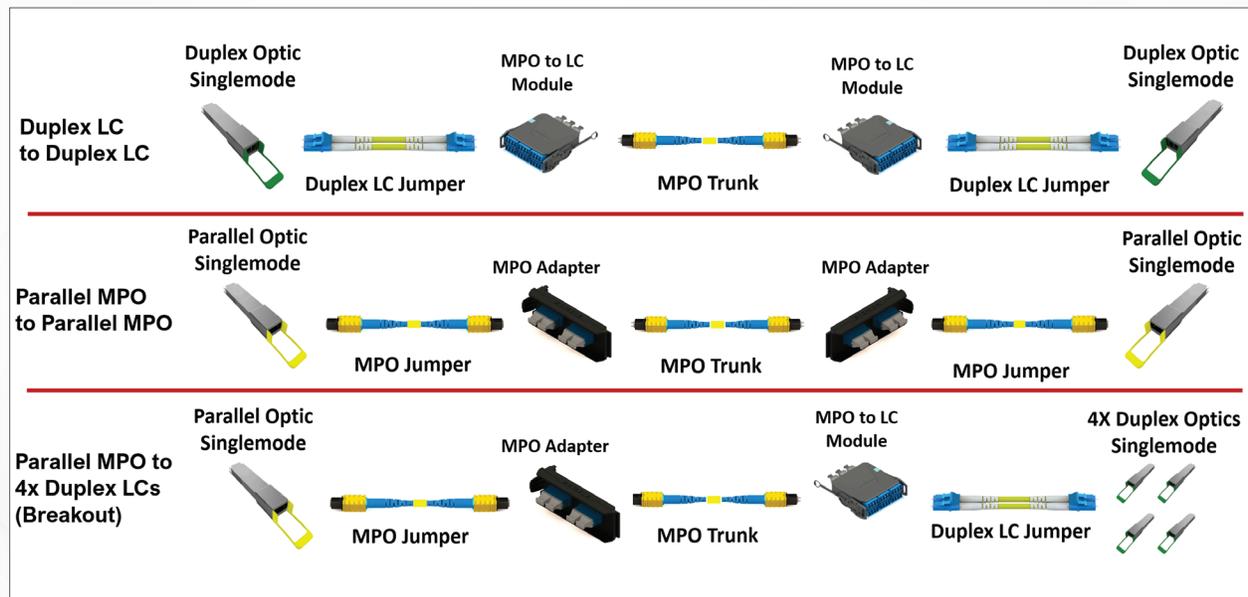
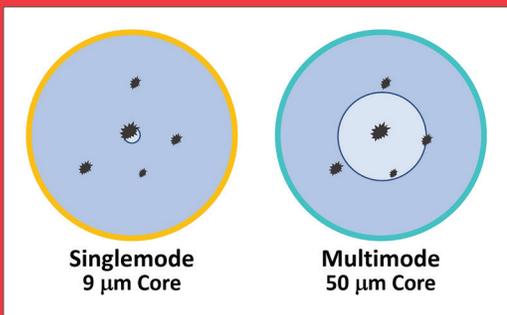


Figure 3: Common data center topologies.

Singlemode Fiber Cleaning Considerations

Fiber end face cleanliness has an impact on both insertion loss and reflectance performance, which is even more critical with singlemode fiber versus multimode due to its smaller core size. A speck of dust on a singlemode fiber end face blocks more of the light signal, and it can even block all of the light signal as shown in this graphic. That's precisely why the IEC 61300-3-35 standard for visual inspection of fiber optic connectors is more stringent for singlemode fiber. For example, a multimode fiber core will pass cleanliness criteria with no limit on scratches equal to or less than 3 microns (μm) in size and up to four defects less than or equal to 5 μm . In contrast, a singlemode fiber core cannot contain any scratches or defects to comply with IEC 61300-3-35. Proper cleaning and inspection are therefore vital in any singlemode fiber deployment. Singlemode APC connectors also require an angled camera inspection tip for proper inspection.



It's important to note that loss values vary from vendor to vendor, with many offering a mix of standard loss, low loss, and ULL connectivity. Table 4 shows varying loss values for common singlemode fiber components from four different vendors. Due to the variability, data center designers need to specify required values

for each component or specify vendors that they know meet those values. They also need to ensure that they are using maximum loss values in their loss budget calculations versus typical loss values that are not guaranteed (see sidebar).

Singlemode Fiber Component Loss Values from Four Vendors				
LC Jumper Connector	Vendor A	Vendor B	Vendor C	Vendor D
Standard	0.30 dB	0.67 dB	N/A	N/A
Low Loss	0.25 dB	0.45 dB	0.75 dB	0.35 dB
Ultra-Low Loss (ULL)	0.20 dB	0.25 dB	0.25 dB	0.25 dB
MPO Trunk Connector	Vendor A	Vendor B	Vendor C	Vendor D
Standard	0.60 dB	0.67 dB	N/A	N/A
Low Loss	N/A	0.45 dB	0.75 dB	0.75 dB
Ultra-Low Loss (ULL)	0.30 dB	0.35 dB	0.35 dB	0.50 dB
MPO-to-LC Module	Vendor A	Vendor B	Vendor C	Vendor D
Standard	1.0 dB	N/A	N/A	0.75 dB
Low Loss	N/A	1.05 dB	1.0 dB	N/A
Ultra-Low Loss (ULL)	0.50 dB	0.60 dB	0.60 dB	N/A

Table 4: Loss values vary across vendors.

Typical vs. Maximum Insertion Loss

Data center designers calculate their insertion loss budgets by adding the insertion loss values for all fiber components in a channel, including the loss of the cable for a given length and the loss of any cable assemblies, jumpers, connectors, and splices within the channel. These insertion loss values are published by component vendors as part of their product specifications. However, some vendors publish both typical and maximum values, while others may only publish typical values or only maximum values. This can cause some confusion for designers in determining which value to use when calculating insertion loss budgets.

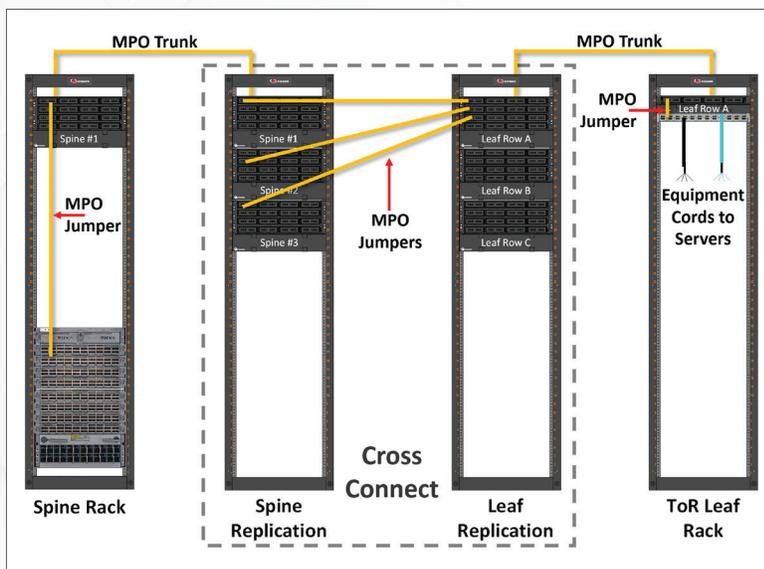
When calculating insertion loss budgets, only maximum insertion loss values should be used. This ensures enough margin for optics to operate correctly and accommodate variables such as poorly installed cabling, dirty end faces, misalignments, aging transmitters, and other factors. In contrast, typical insertion loss values are generated primarily for marketing purposes from median product testing based on typical termination and polishing techniques. Because typical loss values are not guaranteed and may not be a true indication of performance, designers should consider avoiding components from vendors that only publish typical loss values.

Designers should be aware that standard loss connectivity may not support the more stringent insertion loss requirements of short-reach singlemode applications. With each connection in a channel adding insertion loss, designers also need to carefully consider the number of connections and loss values needed to support their preferred data center configuration and provide additional margin to accommodate variables such as poorly installed cabling, dirty end faces, misalignments, aging transmitters, and other factors. For example, the implementation of a cross-connect within a switch-to-switch link adds additional connection points to a two-connector channel, resulting in a four-connector channel.

Even low-loss connectivity can prevent implementing a cross-connect and/or providing sufficient margin. Supporting the number of connections needed to accommodate convenient cross-connects in a channel, while also providing additional margin, can only be guaranteed with ULL connectivity. This is especially the case with MPO-to-LC duplex connectivity commonly used in duplex applications and breakouts since these modules include the loss of both the MPO connector and the LC connector. In Table 5, it is clear that ULL connectivity easily supports the four connections needed to accommodate convenient cross-connects, with plenty of margin to ensure performance and reliability in short-reach singlemode DR and FR applications.

Singlemode Application	Max Loss	Fiber Loss	# of MPO-to-LC Modules			# of MPO Adapters		
			Std 1.0 dB	Low 0.6 dB	ULL 0.5 dB	Std 0.65 dB	Low 0.45 dB	ULL 0.3 dB
DR	3 dB	0.2 dB	2	4	5	4	6	9
FR	4 dB	0.8 dB	3	5	6	4	7	10

Table 5: The number of standard loss, low loss, and ULL MPO-to-LC Modules and MPO adapters allowed in a short-reach singlemode DR and FR channel based on common maximum loss values.



Deploying cross-connects and ensuring sufficient margin is vital for data centers to achieve flexibility, manageability, scalability, speed of deployment, and reliability. In a cross-connect configuration, fiber patch panels replicate switch ports to provide an “any-to-all” configuration, where any switch port can be connected to any other switch port using fiber jumpers at the cross-connect as shown in Figure 4 for a 400GBASE-DR4 spine-leaf deployment.

Figure 4: Cross-connect in a 400GBASE-DR4 channel features fiber patch panels that replicate switch ports, providing an “any-to-all” configuration, where any switch port can be connected to any other switch port using fiber jumpers.

Data centers have long relied on the any-to-all configuration enabled via cross-connects to provide the following benefits:

- **Greater design flexibility** — Cross-connects enable dynamic network design and support creating separate functional areas/server clusters.
- **Easier management and moves, adds, and changes** — Any switch port can be connected to any other switch port using fiber jumpers at the cross-connect. This allows active spine switch ports to be connected to any leaf switch port one at a time to maximize port allocation efficiency and reduce unused active ports.
- **Faster deployments and expansions** — New leaf switches can be easily connected to unused spine switch ports at the cross-connect.
- **Reduced potential for failures** — High-density connections are easily managed at the cross-connect, reducing “spaghetti” cabling that can occur over time and potentially cause failures.
- **Improved security** — In a cross-connect, critical switches can remain untouched and secure, which is especially ideal for data centers that maintain separate switch and server teams and for colocation data centers to establish secure service provider meet-me rooms.
- **Improved traffic flow** — In virtualized and software-defined networking (SDN) environments where resources are often distributed across multiple servers, cross-connects enable easily connecting dispersed resources to the same switch.

The Value of Third-Party Verification

Third-party verification indicates that a component has been certified by an independent accredited third-party body to comply with specific standards for safety, quality, or performance. When products are third-party verified, the accrediting body typically conducts random testing of components at their accredited test lab based on recognized industry standards. These verifications have long played a vital role in instilling consumer confidence across the ICT industry and giving data center owners and operators peace of mind that the components they deploy will indeed comply with industry standards, meet vendor specifications, and ensure the performance they require.

To verify maximum insertion loss specifications and return loss performance for ULL singlemode

connectivity, third-party testing was performed by Intertek, a well-known and trusted inspection, product testing, and certification company in the ICT industry. Testing was performed using eight MPO-mated pairs in a 500 m 400GBASE-DR4 application and with five MPO-to-LC modules and one MPO-mated pair in a 2000 m 400GBASE-FR4 application as shown in Figure 5 (page 11). Note that while channels with this many connections are not common, the testing was conducted to represent a worst-case scenario. The result of the testing clearly shows that ULL, high-performance connectivity provides significant insertion loss and return loss margin for both applications. Remember, the lower the insertion loss value, the better the performance. The higher the return loss value, the better the performance.



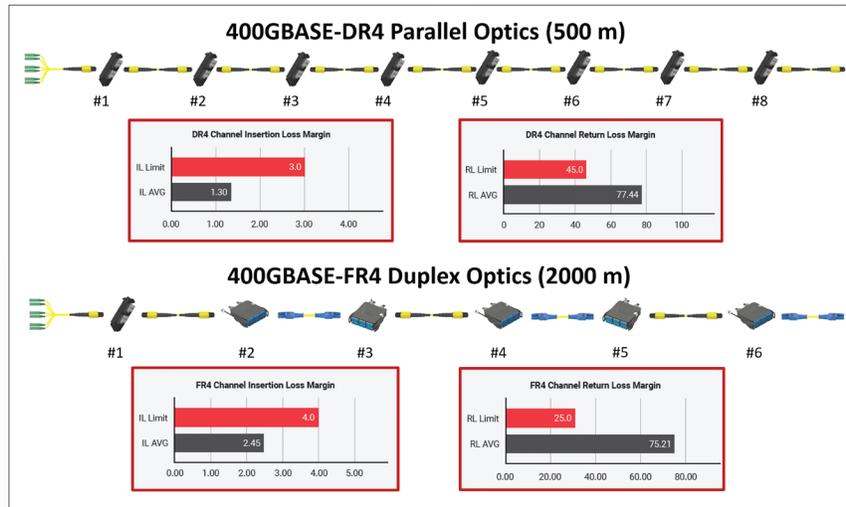


Figure 5: Third-party testing results of ULL, high-performance singlemode connectivity in 400GBASE-DR4 and 400GBASE-FR4 applications.

As shown in Figure 6, testing was also conducted to demonstrate the value of ULL components compared to standard loss components in a cross-connect scenario for both 400GBASE-DR4 and 400GBASE-FR4 applications. In a DR4 channel, the use of standard loss components provides a very minimal margin. Given installation variables, there is the potential for standard loss

components to ultimately exceed the maximum DR channel insertion loss of 3.0 dB, preventing the system from functioning altogether. In FR4 channels that use higher loss MPO-to-LC modules, the use of standard loss components exceeds the maximum channel insertion loss of 4.0 dB and therefore does not support the use of a beneficial cross-connect.

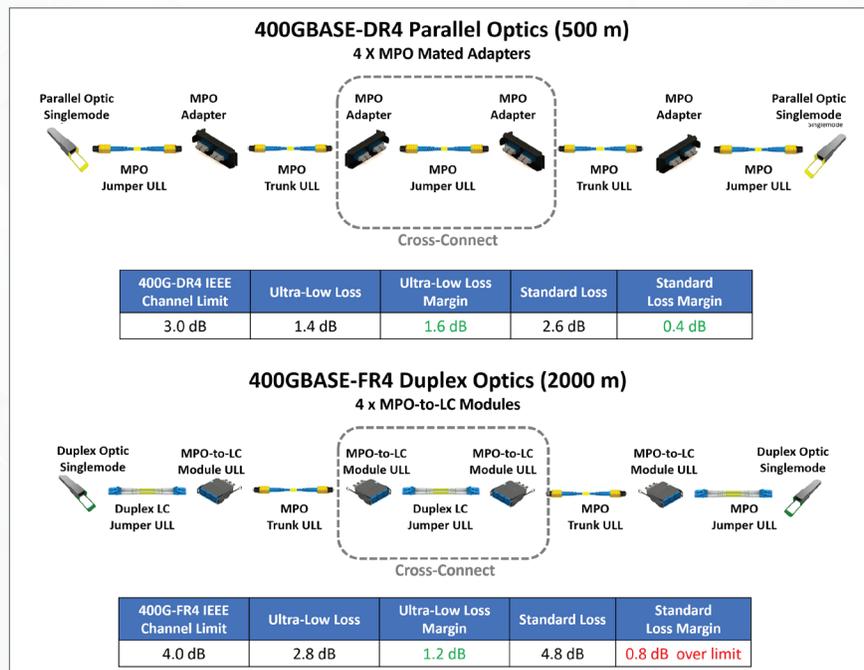


Figure 6: Comparison of ULL and standard loss connectivity in a cross-connect configuration for 400GBASE-DR4 and 400GBASE-FR4 short-reach singlemode applications.

Summary

Enterprise data centers are adopting the hyperscale and cloud data center trend of using singlemode fiber to support higher 100, 200, and 400 Gb/s speeds for emerging applications, while reaching distances beyond the 100 m supported by multimode fiber. Newer short-reach singlemode applications are a far more cost-effective solution for supporting these speeds to greater distances, with DR applications supporting up to 500 m and FR applications supporting up to 2000 m. However, with more stringent maximum channel insertion loss requirements and the need for superior return loss performance, data centers need to pay close attention to vendor insertion loss and return loss specifications.

To enable the flexibility, manageability, scalability, speed of deployment, and reliability that data center managers have valued via cross-connect configurations, while ensuring ample margin to accommodate installation variables, new short-reach singlemode applications need ULL connectivity. In other words, insertion loss matters when selecting components for these cost-effective applications. And for complete peace of mind, data center owners and operators should consider the value of third-party testing verification.



Meet The Author

Gary Bernstein, RCDD, CDCD is the global data center solutions specialist at Siemon. He has more than 25 years of ICT industry experience and extensive knowledge in data center infrastructure, telecommunications, and copper and fiber structured cabling systems. Gary has held positions in engineering, sales, product management, marketing, and corporate management throughout his career. He has been a member of the TIA TR42.7 Copper Cabling Committee, TIA TR42.11 Optical Fiber Committee and various IEEE802.3 task forces and study groups, including 40/100G “ba”, 50/100/200G “cd”, 200/400G “bs” and Beyond 400G. Gary has spoken about data center cabling at several industry events, including AFCOM, BICSI, Cisco Live, and Datacenter Dynamics.

Furthermore, he has authored several articles in industry trade publications. Gary received a Bachelor of Sciences in mechanical engineering from Arizona State University. He is also a BICSI-credentialed RCDD and a Certified Data Center Designer (CDCD) with Datacenter Dynamics. Connect with Gary on LinkedIn. Gary can also be reached at gary_bernstein@siemon.com.



For more information visit:
go.siemon.com/DCSolutions



Find your local Siemon distributor:
go.siemon.com/DCDistributors



24/7 Customer Support:
customer_service@siemon.com

Because we continuously improve our products, Siemon reserves the right to change specifications and availability without prior notice.

North America

P: (1) 860 945 4200

Asia Pacific

P: (61) 2 8977 7500

Latin America

P: (571) 657 1950/51/52

Europe

P: (44) 0 1932 571771

China

P: (86) 215385 0303

India, Middle East & Africa

P: (971) 4 3689743

Siemon Interconnect Solutions

P: (1) 860 945 4213

www.siemon.com/SIS

Mexico

P: (521) 556 387 7708/09/10

WWW.SIEMON.COM

WP_OptimizingShortReachSMDeployments_RevA 10/22

