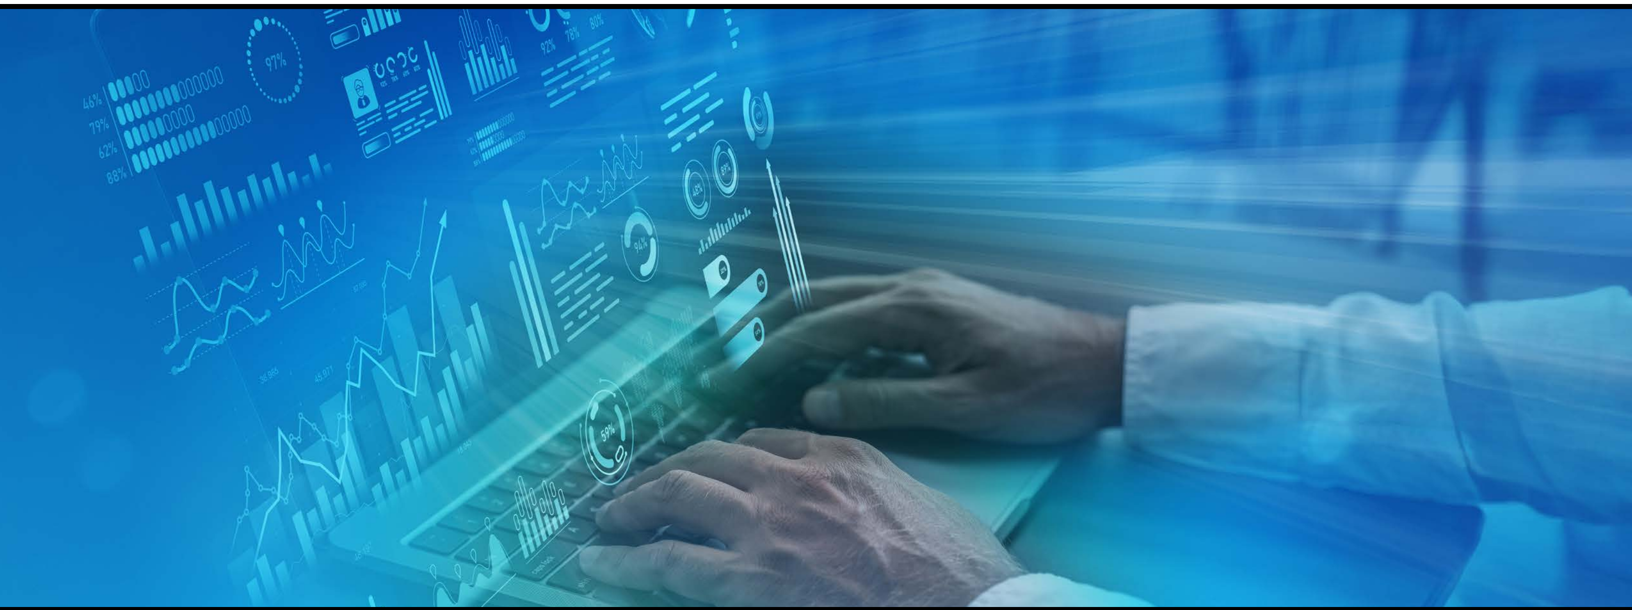

Cabling Strategies to Support Your Data Center's Business Model

There are several schools of thought when it comes to cabling networks, usually based on the type and purpose of the network. Decisions regarding how to cable these networks have changed over the years as networks have evolved. However, the recent rapid expansion of data centers to support AI applications, especially generative AI (GenAI), has significantly changed these models. Today's data center designs are highly dependent on the business model of the data center.

Key considerations include the density and volume of cables, as they have increased by 4x – 8x compared to traditional network cable counts; managing that density to and across cabinet rows; and costs, including not only the materials and labor to install, but also the time, labor and material associated with network life cycle and migration.

Data rates continue to increase, providing greater responsiveness and capacity for users. AI servers or nodes use graphics processing units (GPUs) to deliver greater capacity with each new generation of processor, released with a cadence of every year—or even sooner. Shorter refresh cycles may be justified—if the value outweighs the cost.

Power consumption has grown with newer AI processors, as have equipment cooling demands, requiring innovation in each infrastructure area to support the network ecosystems. Pathways and spaces to manage connectivity within cabinet rows have not significantly grown, but the volume of cables inside and between cabinets has increased by at least 4x – 8x compared to previous networks. This article focuses on this growing cable density, as it directly impacts the flexibility and life cycle of the networks operating across the cabinet row.

With GenAI, there are several cabling options which can coexist within a network. At a high level, the options are either point-to-point or structured cabling, both of which are available in copper and fiber-optic variations. There are benefits and tradeoffs for each, determined by required data rate, distance, flexibility, power efficiency and cost, so the determining factors must depend upon the data center use case or business model.

Cabling options

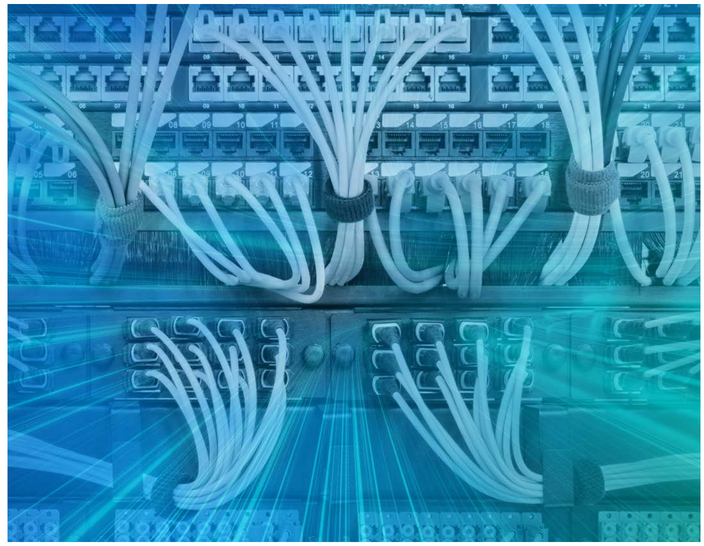

Point-to-point: Point-to-point (PTP) cabling has network connectors attached on each end with built-in transceivers. These cables are single speed, generally protocol-specific and either passive or active, depending on the application. The varieties of PTP are direct attach cable (DAC), which is a passive coaxial cable; active electrical cable (AEC), which is also a coaxial cable but draws power from the electronics; and lastly, active optical cable (AOC), which transmits over fiber-optic cable and draws power at each end as well.

These cables can be one-to-one connections or one-to-multiple for breakout applications. Figure 1 to the right show possible options. The smaller connectors on the ends are SFP+ for applications such as 10G or 25G. The larger connectors represent QSFP, QSFP-DD, or OSFP for applications 50 – 100G and above. The breakout designs can take 400G on one end and break-out to 2x 200G or 4x 100G, for example.

Direct attach cables (DACs) provide a simple point-to-point connection between network devices for short-reach applications. They are considered inexpensive connections within a cabinet, for example, from a top-of-rack (ToR) switch to servers within the rack. These coaxial copper cables are application-specific with SFP or QSFP interfaces at each end and they must be removed and replaced as speeds change. They are typically distance-limited to 1 – 2 m or less; these length limitations get shorter as the data rates increase. DAC cables have larger diameters as lengths increase due to extra shielding required for application support, and they are more rigid with larger bend radii as a result. These cables are passive and do not directly draw power from the electronics but carry 0.15 watts per end from the electronics to transmit signals between devices.

Active electrical cables (AECs) are also single-application coaxial cables with the SFP, QSFP, QSFP-DD or OSFP PC board interfaces on the ends, but, as the name implies, they are active. That means they draw power from the electronics to support higher data rates, but also to drive the signals further. They may draw 0.5 – 1.0W per end, up to approximately 50% of the power needed for transceiver applications end, but they can draw more for higher data rates or additional functionality. They are typically limited to ≤7 m reach and are generally used within or between adjacent cabinets. Similar to DACs, these cables are application- and (typically) protocol-specific, twin-axial copper cables and are thicker and less flexible as data rates or distances increase. Their rigidity can make them impractical for spanning multiple cabinets in a row and a challenge to remove and replace. Considerations for migration to higher speeds should include repurposing or recycling the assemblies.

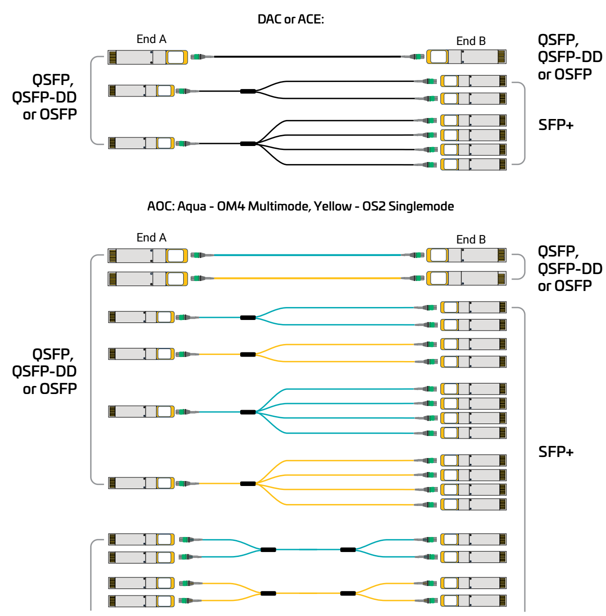

Figure 1. Point-to-point cabling options

Active optics cables (AOCs) are the fiber-optic equivalent of DACs/AECs and are used primarily for higher-speed and longer-reach applications. They have higher data rate capabilities and can be built to support extended distances with QSFP, QSFP-DD or OSFP connectors on each end. As with DACs, they are normally application- and protocol-specific and must be removed and replaced with a network upgrade. They are point-to-point between network devices but utilize smaller cable diameters because of their fiber construction. They can be used within a cabinet, across a cabinet row, or further, to network equipment using singlemode or multimode fiber. As application- and protocol-specific infrastructure, they also must be removed and replaced as upgrades occur. Considerations for their use should include refresh rates and the pathways utilized. Most materials within AOCs are not recyclable, so disposal upon refresh should be properly coordinated. AOCs are used for AI applications at 400G and above. They typically draw 7 – 15W per end, generally in the 12 – 15W range, so proper planning for power and cooling needs is critical.

When running AOCs between cabinets or data center sections it is also important to recognize the weight and bulk of the connector ends which house the transceiver function. As they use small-diameter fiber-optic cables, care must be taken to minimize weight and stress at the connectors. This is especially a concern when entering or exiting cabinets where other cables or obstructions are present. Guidance from leading AI OEMs is to avoid using AOCs for data rates of 100G or higher due to weight and bend stress from the connector at the end of the smaller fiber cable assembly. For higher speeds, structured cabling with transceivers is recommended instead.

Structured cabling systems (SCSs) use passive cable and connectivity and are by nature application- and protocol-independent, supporting multiple network generations over multiple connection points between locations and equipment as needed. For fiber-optic cables, passive connectors such as MPO12/8, MPO16 or LC duplex are terminated on each end for plugging into a transceiver at the network equipment. Fiber-optic transceivers draw similar power from the electronics at each end as AOCs, since they utilize power to convert electrical signals to optical. For applications of 400G and above, they typically draw 7 – 15W per end, and, like AOCs, generally in the 12 – 15W range per end.

With structured cabling, high fiber count, pre-terminated trunk cables are installed between network locations, typically via overhead fiber raceway. Those overhead trunk cables provide connectivity between locations within the data center architecture and also between locations within the cabinet row or POD. Patch panels are utilized within or above the cabinets to provide transition or disconnect points between the trunk and equipment cables. That disconnect point provides flexibility for equipment patching and growth within the cabinet. Fiber patch cables from the panel connect to end devices within the cabinet, completing the circuit from the trunk cables above. If the architecture model can use preconfigured cabinets, then a network upgrade can be quickly accomplished by disconnecting trunk cables from the cabinet, rolling out the original cabinet, and replacing with the upgrade cabinet by reconnecting the trunk cables.

Figure 2 shows is a simple SCS channel using MPO trunk through adapters to breakout cables.

Figure 2. SCS channel using MPO trunk through adapters to breakout cables

Figure 3 depicts a representative baseline AI cabinet row including 32 GPU nodes within the eight outer cabinets and two middle-of-row management cabinets housing aggregated leaf-and-spine switches, compute and storage connections. The cabinet row cable counts are in the lower left of the graphic, recognizing 48 total MPO12/8 connections in each of the eight cabinets.

If cables exiting the cabinets are point-to-point, then the cabling effectively stitches the cabinets together and possibly to the overhead raceway. That may be acceptable for several years of operation in a static network, but when change is needed, those single-application cables must be removed, discarded or recycled and replaced. That will take time, affecting network availability. For dynamic networks like GenAI in particular, flexibility for change may be especially critical.

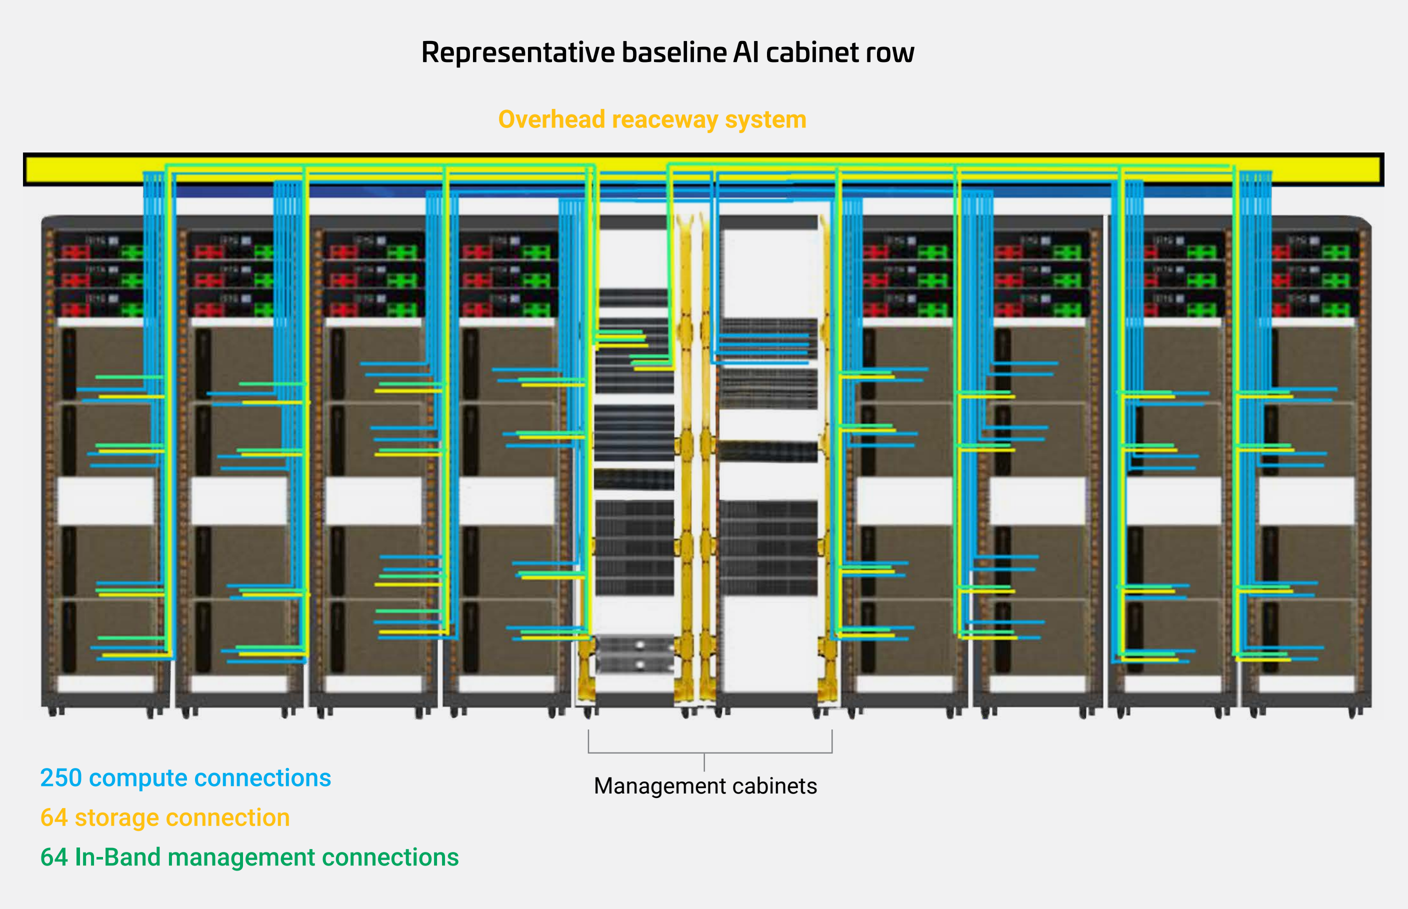

Figure 3. A representative baseline AI cabinet row

It should also be noted that Figure 3 is of an early generation scalable unit, with current options providing over 432 MPO12/8 connections to support 256 GPUs in the space of 5 – 6 cabinets (including switch cabinets).

Relative time, reach and installed cost comparison:

Media	Application	Reach at 100G lanes (400G or 800G ports)	Reach at 200G Lanes (800 or 1.6T ports)	Cost Per Link	Installation Time (1-5)	Applications	Power
DAC (CR)	Ethernet CR	2 m	1m	<x	1	Server to ToR	Low
AOC MMF	Ethernet VR, SR	50 m, 100 m	50 m	2x	4	Server to MoR, EoR	Medium
AOC SMF	Ethernet DR	500 m	500 m	2x	4	T1 to T2	High
Parallel MMF	Ethernet VR, SR	50 m, 100 m	~50 m	x	2	Server to MoR	Medium
Parallel SMF	Ethernet DR	500 m	500 m, 2 km	2x	4	T1 to T2	High
Duplex SMF	Ethernet FR	2 km	500 m, 2 km	>2x	3	T2 to T3	Highest

Conclusion

There are benefits and compromises for each cabling option, based upon data rate, distance, flexibility, power and cost. By design, GenAI networks repeat configurations across cabinet rows and switch layers to build and scale capacity. Those repeated configurations enable the use of pre-configured cabinets built offsite, then rolled in and connected to trunk cables and other infrastructure systems to reduce time to network availability.

Regardless of the data center type, the cabling system used should both support the initial mission of the data center and provide a foundation for future applications and migrations.

By Kenneth Hall, Solutions Architect,
Data Center Solutions

This article was originally published as a chapter in CommScope's 2026 Data Center eBook, available for free here.

EB-117988.3-EN (03/26)