

## Reducing Time to “Go-Live” for GenAI Deployments

**Generative AI (GenAI) data centers need to build quickly to provide availability. Project schedules are getting compressed, from years to months or weeks, and deployed fiber counts have exploded from the thousands to hundreds of thousands.**

Generative AI (GenAI) data centers need to build quickly to provide availability. Project schedules are getting compressed, from years to months or weeks, and deployed fiber counts have exploded from the thousands to hundreds of thousands.


These two evolutionary trends have forced operators to look to their partners for solutions. Certainly, much goes into the planning, construction and delivery of these sites. As complex as these data centers are with their infrastructure power, cooling, and network architecture requirements, their need to scale and provide availability at a rapid pace requires constant deployment innovation. Time efficiencies onsite require a combination of infrastructure readiness, offsite staging and preparation of common components designed and packaged for modularity,

coupled with localized product availability and processes to speed installation and deployment.

An important onsite challenge is managing the flow of the builds. While the scale of GenAI data centers can vary based on the site, the architectures are built upon logical, scalable building blocks to provide capacity across the data center. That modular, repetitive approach—from data halls down to cabinet rows and purpose-built cabinets—provides the foundation for simplifying cabling structure and planning.

### **Point-to-point: advantages and disadvantages**

Consider that compute sections are grouped by cabinet row with common layouts across each multi-GPU compute cabinet. Within each cabinet is a predictable number of cable assemblies between devices in the cabinet; also, those exiting the cabinet to connect to a leaf switch in the management cabinets. As these cable assemblies are consistent throughout the compute cabinet architecture, they can be produced for quality assurance, tested and verified offsite. Patch cords between the top-of-rack panel



and node connections can be installed in cabinets offsite for roll-in and connect network scaling. Trunk cables can be strategically placed overhead in raceways above the cabinet row for quick connections between switch and server cabinets, saving valuable onsite time and resources. The same thought process can be implemented for other cabinet types.

That noted, some instances within a cabinet can utilize point-to-point cabling, such as direct attach cable (DAC), active electrical cable (AEC) and active optical cable (AOC) assemblies which can have initial cost and/or power benefits. As network speeds increase, however, those point-to-point cables will need to be replaced, discarded and upgraded. With copper DAC or AEC assemblies, the rigidity of their coaxial construction may make them challenging to remove, recycle or responsibly discard—only to replace them with more rigid high-speed assemblies.

Point-to-point solutions are even more challenging when the cables need to exit the cabinet for routing elsewhere in the cabinet row. DAC assemblies have short distance limits, and additional shielding for extending this limited reach reduces their flexibility. AOCs are similarly challenging between cabinets, as the weight of high-speed transceivers on the cable connection increases the possibility of damage or breakage of the fiber cable. For these reasons, structured cabling is recommended instead, as it is application- and speed-agnostic and supports multiple network generations.

## Cable construction options

Cable construction is an important consideration. Not only is the routing plan critical, but pathway space for initial installation, along with flexibility for Day 2 additions for migration, must be considered. Individual bundles of 8- or 16-fiber cables in high quantities can quickly fill pathways, whether inside or overhead between cabinets. Should you prefer the granularity of individual cords, consider using 16-fiber assemblies, since they occupy less volume than comparable fiber counts with 8 fiber units. A more efficient cable footprint would come from higher fiber count assemblies with smaller bundle size with 8- or 16-fiber breakouts at appropriate locations, using MPO connectors to connect at the transceiver.

New smaller very small form factor (VSFF) multipair connectors are in early market deployments, which can initially bring reduced-size benefits for trunk and pathway applications. MMC16 or MMC24 and SN-MT16 connectors deliver the same fiber count as their traditional MPO16 or MPO24 counterparts in roughly one-third the space. Their smaller construction significantly reduces the area required in the vertical or horizontal pathway, making it easier to pull assemblies between locations in the data center.

The market is developing for these connectors, as well as for emerging ganged connectors which also leverage VSFF constructions to provide high fiber counts in a much smaller connector footprint. The expectation is that they will reduce the time needed to connect the fibers and reduce the risk of connectors being installed into the wrong port.

## Connections in evolving data center environments

Depending on the architecture, cable connections can occur on the front or rear of a cabinet (or both), based on their purpose. They may need to be routed around power devices or cooling hoses. When possible, plan to use larger or taller cabinets to maximize space for the systems installed.

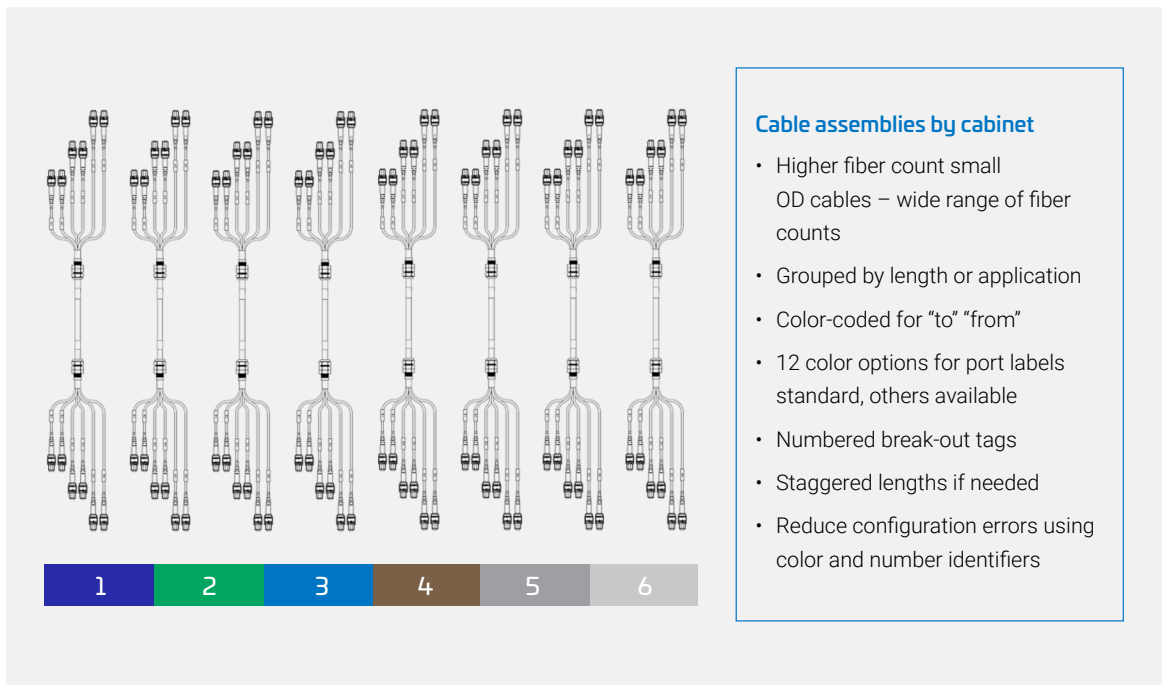
Recent deployments scaling up and out of AI data center PODs have brought to light the challenges of collocated systems in cabinets, mainly GPU compute cabinets with limited vertical space on the front, along with CPUs for management, hoses for direct to chip (DTC) liquid cooling and power. These cabinets are generally 4 feet deep or deeper and, as such, are problematic to cable using traditional structured cabling methods.

Obstructions within the cabinets are only a part of the challenge. Human arms are generally shorter than 4 feet long, so accessing the panel connections from the rear is not practical. Solutions that are front-facing-only provide possible options using high-density VSFF connections providing the depth and height are factored into the design. In general, panel space is limited to 1 – 2 RU in a GPU cabinet due to the presence of other network devices necessary for functionality.

Bear in mind that direct connecting GPU cabinet transceiver ports to leaf switch ports can be accomplished as front-facing assemblies by using cable bundles. Managing the identification, routing, and protection of slack and spare cables, along with the extended breakout lengths over and between cabinets, is challenging for Day 1 installers. Day 2 becomes harder to access or change as the cabinets get stitched together with that point-to-point design. Running individual patch cables has challenges as well; cable bulk from the volume of several hundred 8- or 16-fiber cords between cabinets is difficult for installation as well as troubleshooting.

Accommodating and routing cable through available pathways relies on a well-planned approach. Pre-measuring known cable paths—building in a manageable slack factor—is important for installation and change management. Once length is verified, one must determine timing and volumes required. It's important to plan this as far in advance of deployment as possible, as custom assemblies purchased in volume may require additional lead time.

In addition, as GenAI deployments require groups of MPO assemblies at each GPU to deliver higher fiber count connections back to the management cabinets, vertical and horizontal pathways can quickly become congested. Planning for that cable mass is critical to Day 2 operations when cables may need to be added for next-generation applications. In a short period of time, we have seen the migration from 400G using eight fibers to 800G using 16 fibers (2x8f); soon, we'll see 1.6 Tbps becoming mainstream for the larger AI sites, possibly doubling the fiber count again. For backward compatibility, using 100G lanes would double the fiber count. With 200G lanes on the horizon, there will be options to maintain the current fiber counts with 2x speed. Either way, planning cable assembly configurations can simplify installation and enable flexibility for operations teams.



For example, within spine-leaf switch and management cabinets, high fiber count array assemblies color-coded with numbered labels can simplify and speed up onsite installation. Since these assemblies are replicated through each cabinet row, they can be pre-measured for pathways and placement. These can be made in various port and fiber counts and break-out lengths to provide efficiencies upon installation. Assemblies with staggered lengths can simplify proper placement.

To save time onsite, these assemblies—along with the required network equipment—can be pre-installed in a cabinet offsite by a value-add partner. Offsite or off-data-center-floor preparation can reduce effort expended by simply dropping trunk cables from overhead into the top of the rolled-in cabinets and connecting to the modules at the top-of-rack fiber panel.

If fit-out were done onsite, it would require unpackaging, cleaning and inspecting the assemblies, determining which assemblies are positioned where within the cabinet, neatly routing and managing the cables in the pathways (ensuring to and from connections are complete for dozens of cables in each cabinet), and then connecting trunk cables to the panel at the top of the cabinet.

Advance steps can be taken to improve efficiency when materials get to the site. From a fiber connectivity perspective, overhead pathways should be installed before the cabinets are brought into the data center. Factory-terminated and tested trunk cables should be provisioned above the planned cabinet locations. Labeled trunk cables provide fiber routing between endpoints in the cabinet row or across the data center. All fiber patching can be completed within this cabinet to the front of a fiber panel at the top of the cabinet. Once arriving onsite, the cabinet is positioned, trunk cables are dropped and routed into the cabinet and are then connected to the installed equipment cords via the fiber panel.

Aside from simplifying and speeding up the installation process, the building block approach provides the foundation for future rapid migrations via disconnecting the cabinet trunks, rolling the cabinet out for next steps and then rolling in and reconnecting trunks to the next-generation cabinet.

## Steps to simplify deployment

Whether the cabinets are configured and cabled offsite or onsite, cable assemblies can be built, organized and packaged to simplify the process and reduce time required onsite. For example, all of



the cable assemblies for a single cabinet could be factory built, tested, color-coded and number labeled, packaged, and marked with “to” and “from” labels for each of the requirements in that cabinet. QR-coded assemblies can provide all factory test data and documentation for each assembly.

Another option would be to package up all similar purpose assemblies for the cabinet row, with similar labeling. Examples would be node-to-fiber panel at top of the cabinet, or spine-to-leaf switch connections in the management cabinet. Installers onsite could plug in the assemblies as they work down the row. The idea is to set the workflow and follow the simplified sequences to be more efficient onsite.

## Taking a long-term view

Consider the need to have redundancy and error-free transmissions to reduce the chance that the network fails. In recent years, “mesh” or “shuffle” network cabling provided a way to keep the network operating when a port or switch failed on a duplex fiber level. This is more critical for multi-pair applications.

For example, 400G applications would split into 4x100G ports along with other similar ports in the network. Should a 400G device fail, a small, connected portion of the network would continue operating at a slightly reduced speed (300G), but the network would not go down. Service could be scheduled but would not require an immediate truck roll. GenAI cluster uptime requirements are better supported by incorporating a shuffle element in some critical areas. As most of these cabinet row connections use parallel pairs of fibers for speeds of 200G, 400G and above, redundancy can be included using common configurations. The higher the pair count in connections, the greater the benefit of shuffle design.

## Making the best choice

There are benefits and tradeoffs for each cabling option, with dependencies 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 to be 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(s) used should support the initial mission of the data center with a vision 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](#).

CommScope pushes the boundaries of communications technology with game-changing ideas and ground-breaking discoveries that spark profound human achievement. We collaborate with our customers and partners to design, create and build the world's most advanced networks. It is our passion and commitment to identify the next opportunity and realize a better tomorrow. Discover more at [commscope.com](https://commscope.com)

**COMMSCOPE**<sup>®</sup>  
an Amphenol company

[commscope.com](https://commscope.com)

Visit our website or contact your local CommScope representative for more information.

© 2026 CommScope Technologies LLC, an Amphenol company. All rights reserved. CommScope and the CommScope logo are registered trademarks of CommScope and/or its affiliates in the U.S. and other countries. For additional trademark information see <https://www.commscope.com/trademarks>. All product names, trademarks and registered trademarks are property of their respective owners.

EB-117988.3-EN (03/26)