

# Multicore Fiber for AI-Scale Data Halls: A Lower-Carbon Path to Density, Speed, and Scalability

Aislin K Sullivan  
Shane R Rainey

## Abstract

Data center operators face increasing pressure to rapidly scale AI-ready infrastructure while navigating a critical skilled-labor shortage (Smith, 2026) and decarbonizing operations to meet aggressive greenhouse gas (GHG) reduction targets. This paper evaluates high-density, preterminated passive optical infrastructure – specifically 4-core multicore fiber (MCF) vs. single-core fiber (SCF) – using a cradle-to-gate Product Carbon Footprint (PCF) screening of a representative generative AI data hall architecture (Herrmann, Diaz, Wu, & Von Zimmerman, Mar 2026). The study quantifies carbon footprint across optical fiber, optical fiber cables, and the associated passive optical components that materially differ between two data hall configurations, each supporting a fixed functional compute footprint of 18,432 GPUs. Results show that increased pathway density, enabled by multicore fiber, reduces material demand (e.g., modular housings, cabling, and ducting) and delivers up to 60% lower carbon footprint impact per GPU for the passive optical components in scope, reducing related data hall carbon footprint from 4,549 metric tons CO<sub>2</sub>e (SCF) to 1,792 metric tons CO<sub>2</sub>e (MCF).

Not quantified in this study are the active layer (e.g., transceivers) and the broader facility-level pathway changes (e.g., installation/buildout impacts). While the passive MCF connectivity clearly delivers sustainability benefits—such as space reduction and lower raw material usage—the optical transceiver itself consumes the same amount of power as a standard non-MCF transceiver. Therefore, the passive layer efficiency gains presented here are additive, and are not expected to be offset by, active layer performance. From an embodied carbon perspective, fan-in/fan-out (FIFO) components—implemented as fiber bundles or 3D waveguides in MCF transceivers – are relatively small, passive parts made from common non-critical materials using established manufacturing processes; accordingly, we expect their contribution to total carbon footprint to be negligible.

Finally, although facility-level pathway changes are not quantified in this study, higher density, preterminated solutions directly address the primary physical and economic buildout constraints in AI data halls: skilled labor availability, cable tray congestion, and rack space limitations. As GPU generations advance, and fiber requirements progressively double to support higher bandwidth data transmission, high density, preterminated solutions become a critical enabler of scale, allowing operators to increase connectivity capacity without expanding tray and rack footprint, all while enabling deployment with the existing labor pool.

## Introduction

Generative AI is accelerating the expansion of hyperscale data centers and materially increasing the resource intensity of each new build—driving demand not only for power and cooling capacity, but also for the construction materials and electrical/mechanical infrastructure required to deliver higher compute density. At the same time, the pace of deployment is increasingly constrained by shortages in skilled labor required to build data center infrastructure. Optical connectivity is foundational to this growth by enabling high-bandwidth, low-latency communication between compute nodes.

Although operational electricity used for power and cooling remains the largest lever to reduce resource intensity, there is significant opportunity to optimize the infrastructure itself. As AI fabrics scale, physical constraints in pathways (e.g., cabling routes

and in-rack capacity) increasingly shape cost, schedule, and design feasibility—making higher density optical architectures a practical lever to reduce material use and associated carbon footprint as deployments grow.

This paper focuses specifically on the cradle-to-gate carbon footprint comparison of two data hall configurations that support a fixed functional compute footprint of 18,432 GPUs: one with traditional single-core optical fiber (SCF), and a high-density data hall design enabled by 4-core multicore fiber (MCF). The scope of the study is focused on selected passive optical components that materially vary between each configuration.

### MCF enables density improvements up to 4x

MCF is an optical fiber engineered to carry multiple independent optical pathways (“cores”) within a single fiber, enabling a step change in connectivity density without changing the overall fiber form factor. Corning’s 4-core MCF delivers up to four times the capacity of single-core fiber (SCF) within the same 125  $\mu\text{m}$  cladding diameter and 190  $\mu\text{m}$  coating diameter, while maintaining optical performance. While MCF can have increased microbending sensitivity relative to single-core fiber with an identical refractive index profile, Corning’s MCF is designed with consideration to mode field diameter and coating properties to be a drop-in replacement for today’s single-core fibers in telecommunication cables.

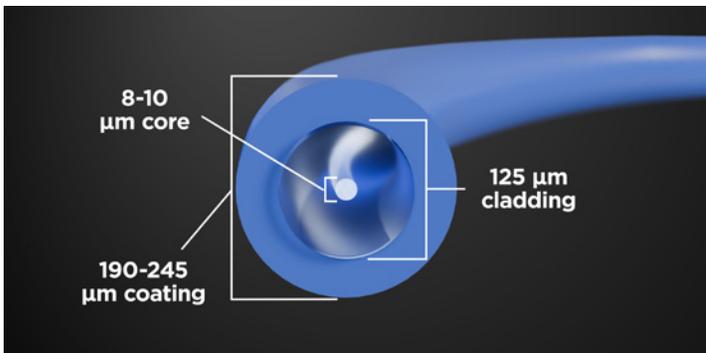


Figure 1: Standard single-core fiber

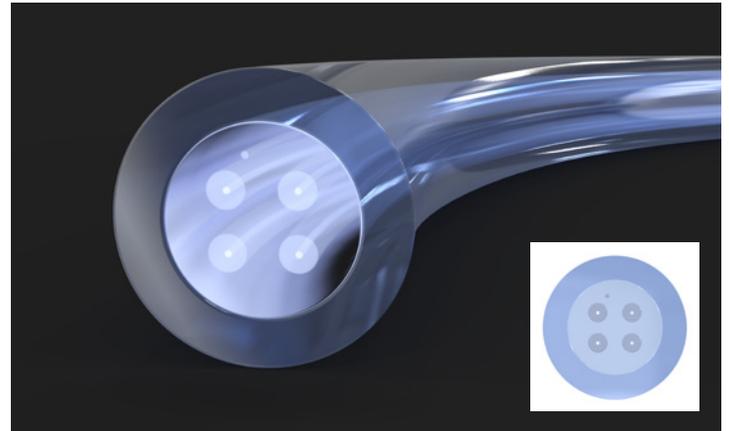


Figure 2. 4-core MCF

### MCF makes higher cable density possible

Loose tube and ribbon are the two dominant cable designs used to package fibers, and the choice materially affects cable density (measured as fibers/ $\text{mm}^2$  or cores/ $\text{mm}^2$ ).

In a loose tube design, fibers are grouped inside small protective buffer tubes and stranded together; this construction is robust and widely used, but it introduces “packing overhead” (air space and tube structures), so cable outside diameter tends to grow faster as fiber counts increase.

Organizing fibers into ribbon structures improves packing efficiency and enables more fibers in a smaller cable diameter. GPU generations are driving steep growth in required interconnect fiber counts well beyond the range where loose tube cables can comfortably fit in existing pathways. Smaller, denser ribbon cables therefore provide an advantage.

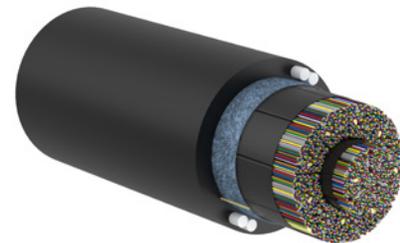
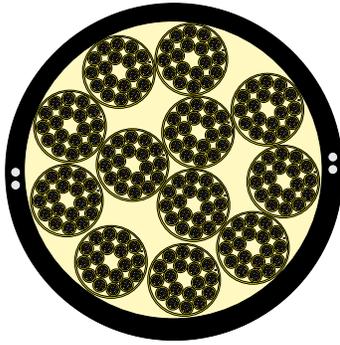


Figure 3: Corning® RocketRibbon® XD Cable-200 Flow, outside plant (OSP), with 3456 SCF delivers high fiber density of 6.5 fibers (cores)/ $\text{mm}^2$



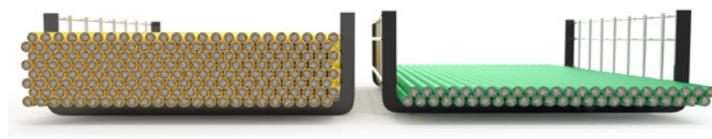
**Figure 4:** 3456F super 16, indoor/outdoor (I/O) riser (loose tube design) with MCF. Density: 3.7 cores/mm<sup>2</sup>

Whether implemented in ribbon or loose tube cable design, the density benefits of MCF become most meaningful when it can be a drop-in replacement in existing cable designs with minimal to no changes to diameter. With that premise, the density advantage is straightforward: independent of whether fibers are packaged as ribbons or placed in loose tubes, carrying 4x more optical cores per fiber increases total optical lane capacity accordingly, helping relieve pathway constraints while preserving compatibility with established cable designs and deployment practices.

Cable Design	Outer Diameter	Total Optical Cores	Core Density
3456F Ribbon Flow with SCF	26 mm	3456	6.5 cores/mm <sup>2</sup>
3456F Ribbon Flow with MCF		3456 x 4 = 13,824	26 cores/mm <sup>2</sup> (4x SCF)
3456F Loose Tube with SCF	34.4 mm	3456	3.7 cores/mm <sup>2</sup>
3456F Loose Tube with MCF		3456 x 4 = 13,824	14.9 cores/mm <sup>2</sup> (4x SCF)

**Table 1:** Core density improvements across ribbon (OSP) and loose tube (I/O) cable design

From a data center operator perspective, the practical value of MCF is that it increases pathway capacity without proportionally increasing cable bulk, helping relieve physical congestion in overhead pathways and cable trays as AI clusters scale.



**Figure 5:** SCF cable (left) and 4-core MCF cable (right) equivalent pathway deployment in cable tray

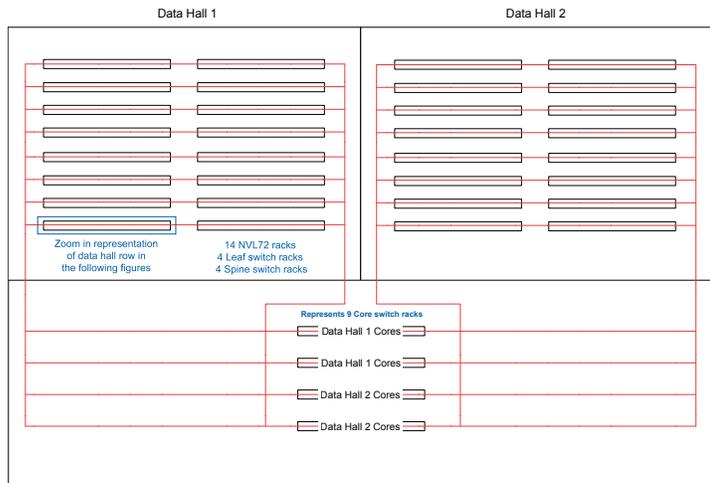
More importantly, MCF solutions simplify network architecture, thereby simplifying passive optical network equipment installation. Corning’s conservative modeling utilizing NECA Manual of Labor Units suggests ~60% reduction in passive optical infrastructure installation and testing labor, and up to a nine-month reduction in the physical-layer cabling deployment and testing schedule for the representative Blackwell network design. The speed of deployment improvement helps alleviate the labor constraint seen by the data center industry.

### Comparing embodied carbon of two data halls

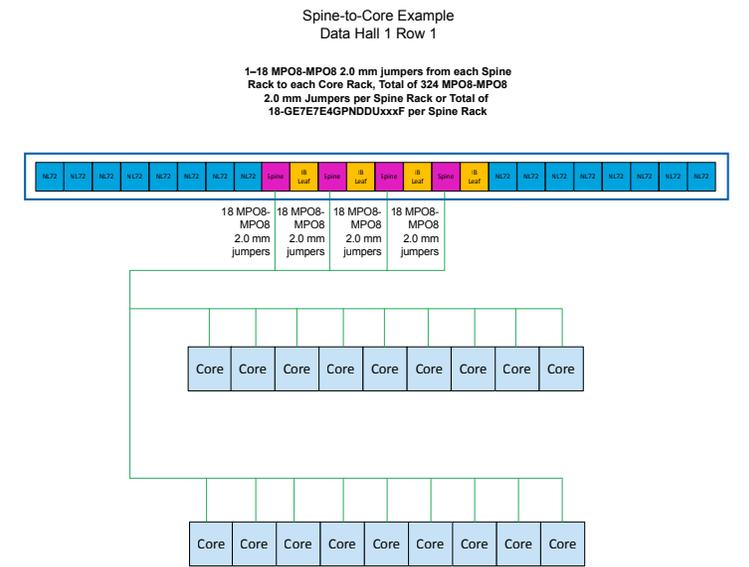
In generative AI environments, data halls are a purpose-built environment designed for extreme computational demands. They typically comprise rows of high-density, GPU-accelerated servers interconnected through high-bandwidth, low-latency networking fabrics to enable large-scale model training and inference.

To sustain continuous operation, the data hall is supported by robust thermal management and redundant power infrastructure. Within this environment, passive optical infrastructure—fiber optic cabling and associated routing and management components—provides the physical transmission pathways that connect compute, switching, and aggregation layers across the network. In practice, this includes cable assemblies (e.g., jumpers and trunks), modular interconnection hardware (e.g., housings), and routing elements such as cable trays and ducts that organize, protect, and scale fiber pathways throughout the hall and between halls.

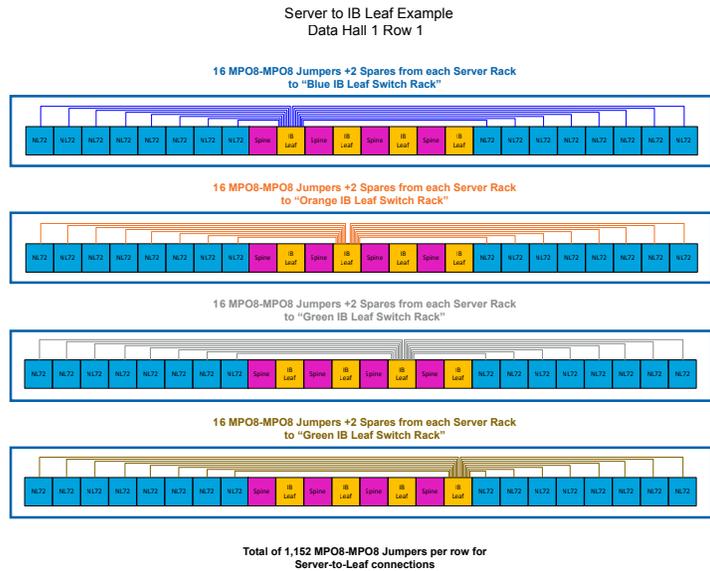
To help interpret the material quantities modeled in each configuration, Figures 6–9 visualize the representative generative AI data hall architecture used. The modeled data hall is constructed as repeated rows of high-density, GPU-accelerated compute interconnected through passive optical infrastructure. The architecture shown was developed by translating a Blackwell-class, rack-scale deployment baseline (modeled around NVIDIA GB200 NVL72 systems and direct-reach pluggable optical connectivity) into a modular leaf-spine-core topology consistent with commonly deployed customer/reference patterns and published scalable AI cluster guidance. Each link class was then mapped into structured passive cabling and pathways so the resulting quantities could be compared for SCF vs. 4-core MCF while holding the data hall function constant at 18,432 GPUs with a 1:1 oversubscription ratio. Figure 6 shows the full row-based design and the associated quantities of network components per row, while Figures 7–9 successively zoom into the passive optical connections from server-to-leaf, leaf-to-spine, and spine-to-core to clarify how the connectivity map drives the modeled cabling and pathway requirements.



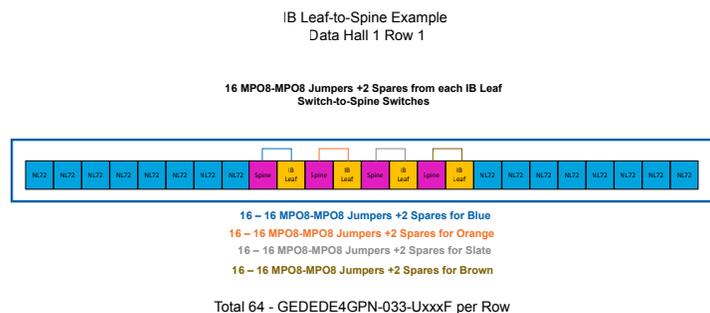
**Figure 6:** High-level AI data hall network design and associated quantity of network components per row



**Figure 9:** Zoomed in view from Figure 6, with inclusion of the core switches to demonstrate the Spine-to-Core passive optical network connections



**Figure 7:** Zoomed-in view from Figure 6 to demonstrate the Server-to-Leaf passive optical network connections



**Figure 8:** Zoomed in view from Figure 6 to demonstrate the Leaf-to-Spine passive optical network connections

In the SCF configuration, each fiber provides one optical pathway, which drives higher total volumes of cabling and pathway materials to achieve the required number of optical lanes. Accordingly, the SCF data hall includes 1,750 km of intra-hall cable (12F MIC<sup>®</sup> loose tube) used in the MPO8-MPO8 jumpers, 268 km of backbone cable (3456F RocketRibbon<sup>®</sup> Flow) used in preterminated EDGE<sup>™</sup> Rapid Connect solution, 134 modular housings (EDGE8<sup>®</sup> 4U), and associated routing infrastructure, including 0.608 km of cable tray and 488 km of ducting (268 km of 1.25-inch duct, plus 90 km each of 4-inch duct and 3-way duct).

In the MCF configuration, each fiber provides four optical pathways, enabling comparable lane capacity with fewer physical cables and reduced supporting infrastructure. The modeled MCF data hall therefore uses 875 km of intra-hall cable (4F MIC loose tube [MCF]) used in the MPO4-MPO4 1.6 mm jumpers, 68 km of backbone cable (3456F ribbon flow [MCF]) used in preterminated solution, 67 modular housings (EDGE8 4U), and reduced routing materials, including 0.456 km of cable tray and 116 km of ducting (68 km of 1.25-inch duct and 24 km each of 4-inch duct and 3-way duct).

## Methodology: A Life Cycle Assessment (LCA) approach

This paper evaluates the embodied greenhouse gas (GHG) emissions, reported as product carbon footprint in kg or metric tons CO<sub>2</sub>e, for selected passive optical infrastructure used in a generative AI data hall. “Passive” refers to optical

components such as optical fiber, optical fiber cables, modular housings, ducts, and trays, and excludes active equipment such as servers, switches, transceivers, and cooling systems.

Corning Optical Communications LLC commissioned Sphera Solutions, Inc. to conduct a cradle-to-gate product carbon footprint (PCF) screening of these components, prepared in accordance with ISO 14067 (Herrmann, Diaz, Wu, & Von Zimmerman, Mar 2026). This considers all the resources needed and emissions caused by raw material extraction (“cradle”), their processing into semi-finished products, and final manufacturing into components that are packaged and ready for further distribution (“gate”).

The resulting LCA report was critically reviewed by an independent third-party and is intended to support internal and external stakeholders with an objective basis for reporting and for identifying impact reduction opportunities.

Corning collected primary data from its production sites and detailed material information to quantify the cradle-to-gate impact for the optical fiber, cables, and housings and background datasets were sourced from Sphera’s Managed LCA Content (MLC) database (Sphera, 2026). The routing accessories were modeled by Sphera based on relevant specifications.

### The screening covers three layers within a data hall:

- Optical fiber (single-core and 4-core multicore)
- Optical fiber cables (including single-core and 4-core multicore cable configurations)
- Other data hall passive optical infrastructure, limited to the passive components expected to dominate cradle-to-gate GHG contributions within scope (modular housings, ducts, and trays).

To ensure equivalent comparisons across the three layers, the study normalizes results to a consistent functional reference for each layer of the system. For optical fiber, impacts are reported per kilometer of fiber on a defined number of optical pathways (one pathway for SCF and four pathways for MCF); for optical fiber cables, impacts are reported per kilometer of cable and expressed on an equivalent “per optical lane” basis, so different cable constructions can be compared on the same delivered capacity. For the data hall, the comparison holds the end function constant by modeling the passive optical infrastructure required to support the same AI data hall scale—256 NVIDIA GB200 NVL72 systems totaling 18,432 GPUs—and then quantifies the materials and components needed to deliver that connectivity.

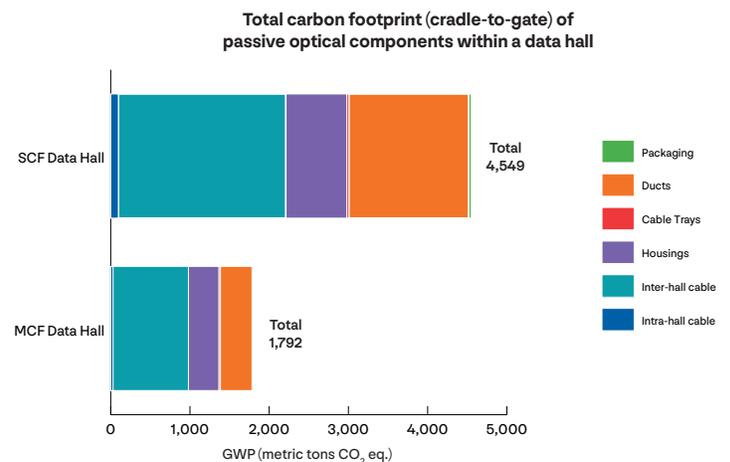
The study excludes active electronics—including transceivers, additional constant passive components treated as

unchanged across configurations (e.g., patch hardware), installation, building infrastructure, transport of components to the data hall, the use stage, and end-of-life of installed systems. For cable assemblies, only the cable portion is included; having minimal impact in the total carbon footprint, terminations/connectors are excluded at this screening stage.

## Data hall results

Both data hall configurations support 18,432 GPUs, but MCF’s higher core density allows for a more compact physical footprint. When comparing the passive optical components in scope, 4-core MCF enables a reduction of >2,700 metric tons of CO<sub>2</sub>e, equivalent to the annual electricity use of 575 homes (Greenhouse Gas Equivalencies Calculator, 2026).

In both cases, the total carbon footprint is largely driven by the contribution of backbone cable and ducts for inter-hall connections, indicating opportunities for further innovation in this space.



**Figure 10:** Carbon footprint comparison of two data halls, one with SCF solutions and one with 4-core MCF solutions (passive optical components in scope only)

For each GPU, 4-core MCF reduces the carbon footprint of the relevant passive optical components by up to 60%.

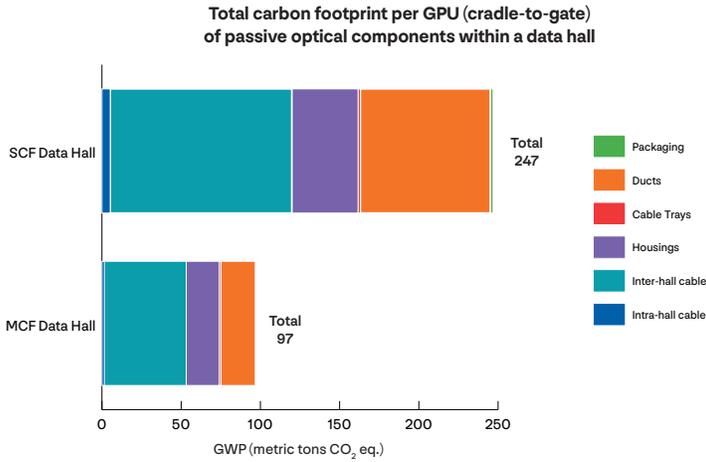


Figure 11: Carbon footprint per GPU, passive optical components in scope only

Often overlooked, packaging represents another key opportunity for driving sustainability. Beyond simply reducing waste, optimized packaging can reduce installation time and accelerate deployments, driving efficiency from end-to-end. Switching from a SCF to an MCF data hall substantially reduces the amount of packaging required during deployment. Total packaging mass drops from 74 to 19 metric tons (55 metric tons or ~74% reduction). The reduction is driven primarily by steel reels, which fall from **134 to 34 reels and from 74 to 19 metric tons of packaging mass (~55 metric tons reduction)**. Additional reductions occur across other packaging types: plastic bags/foam decrease from 76,032 to 47,232 units and from 517 kg to 321 kg, cardboard boxes from 855 to 315 units and from 221 kg to 81 kg, and wood pallets from 11 to 4 pallets and from 165 kg to 60 kg.

### How MCF is made, and why it matters for carbon footprint results

Optical fiber manufacturing begins with high-purity raw materials that undergo a high-temperature vapor deposition process to form silica “soot” with the required optical properties. The soot is deposited onto starter rods and consolidated into a dense glass preform. For SCF, the preform contains one centrally located core. For 4-core MCF, the preform incorporates a more complex structure that requires drilling and the precise placement/fusing of four cores within a single cladding. This added complexity primarily affects the preform fabrication stage and is the primary reason multicore manufacturing is modeled with higher energy demand than single-core manufacturing.

Once the preform is complete (SCF or MCF), it is transferred to a fiber draw tower, where the preform tip is heated and drawn into a continuous glass fiber at the target diameter. The fiber is then immediately coated with protective acrylate layers, UV-cured, and spooled for subsequent cable assembly operations.

As MCF is not yet manufactured at the same industrial scale as SCF, the study treats electricity use in multicore manufacturing as a key sensitivity driver. To remain conservative, the base modeling approach represents multicore manufacturing using a premium/specialty fiber process (lower volume and higher mix) and explicitly tests how results vary with different electricity consumption assumptions. Importantly, even with the conservative assumption of higher electricity intensity per kilometer, Corning’s 4-core MCF design delivers four optical pathways per fiber—so when results are expressed on an equivalent “per core” basis, the carbon footprint advantage is maintained, yielding a 50% lower carbon footprint per optical core.

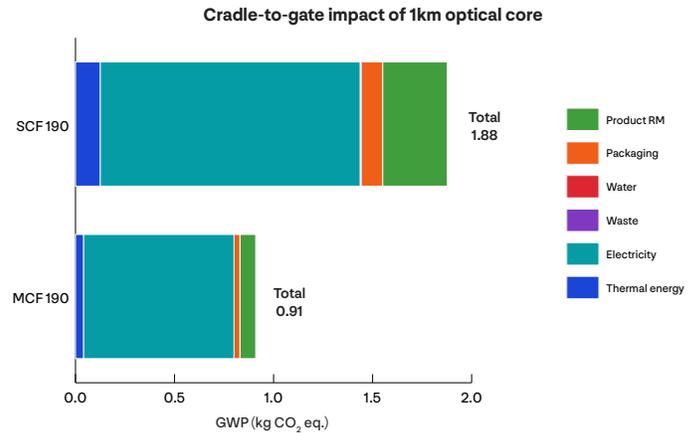


Figure 12: Carbon footprint impact per km per core of single core and 4-core multicore fiber

At the cable level, increased core density means less raw material and a lower carbon footprint per core. 4-core MCF cable can cut the carbon footprint per core by up to 56% compared to single-core designs, and up to 75% when this advantage is compounded with high-density designs.

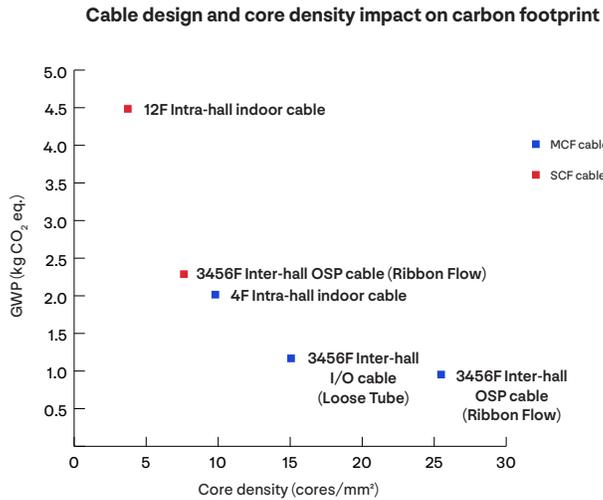


Figure 13: Carbon footprint impact per km\*core based on density improvements (measured as cores/mm<sup>2</sup>)

### Practical recommendations for data center operators

From a data center operator perspective, the practical value of MCF is that it increases optical pathway capacity without proportionally increasing cable bulk, helping relieve congestion in overhead pathways, racks, and ducts as AI fabrics scale. In the cradle-to-gate screening PCF summarized in this paper, 4-core multicore configurations reduce carbon footprint for the passive optical components in scope by up to 60% per GPU, driven primarily by reduced material and electricity demand across fiber, cable, and supporting pathway infrastructure required to deliver equivalent connectivity.

#### Adopt a pathway-constrained design approach early

AI deployments increasingly hit physical pathway limits—overhead tray congestion, rack unit patching space, and finite duct capacity—before they run out of optical performance margin, and 4-core MCF addresses this by delivering up to 4x capacity per fiber without proportionally increasing cable bulk. In the screened data hall comparison, that density translates into fewer required passive components to achieve the same functional outcome, which drives the cradle-to-gate carbon footprint reduction for the passive components in scope.

### Pair MCF with high-density connectivity to compound the space benefit

MCF density benefits are most meaningful when combined with connectivity that preserves rack-level packing efficiency. When integrated with high-density connectors like MMC, high port density is enabled while maintaining practical deployment workflows.

### Use preterminated assemblies and fewer connections to address labor and schedule constraints

Beyond carbon footprint, one of the strongest operator levers is reducing schedule risk. Modeling, based on the referenced architecture for this study, indicates ~60% reduction in passive optical infrastructure installation and testing labor and up to a nine-month reduction in cabling deployment and testing schedule for a representative AI network design, consistent with the “fewer cables, fewer connections” mechanism enabled by higher density.

### Packaging collaboration between data center operators and equipment suppliers can unlock sustainability wins

At the scale of AI data halls, packaging stops being incidental and becomes a real material stream: dozens of heavy reels and large volumes of secondary packaging consume floor space, add handling steps, and create a non-trivial waste and reverse-logistics burden during buildout. The multicore configuration materially reduces packaging mass, which can translate into less staging space, fewer disposal operations, and simplified on-site logistics during rapid buildouts, reducing deployment waste streams.

As AI data halls scale, the limiting factors are increasingly physical: pathway congestion, rack-space constraints, and the availability of skilled labor to deploy complex cabling at speed. This study shows that multicore fiber-enabled, high-density architectures can materially reduce carbon footprint for passive optical infrastructure while alleviating the real-world constraints that slow deployment. For operators seeking to expand AI capacity without proportionally expanding material use, footprint, or labor exposure, pathway density is an actionable lever—and MCF is a practical means to achieve it.

## Acknowledgments

This white paper reflects the contributions of multiple teams across Corning Optical Communications business units. The authors thank Khalid Amao, Rodrigo Martinez Elizondo, and Carly Gaj (Corning Optical Communications) and Constantin Herrmann, Amalia Diaz, Roberta Von Zimmerman, and Tammy Wu (Sphera Solutions Inc.) for their substantial contributions of time, models, and expertise for the Screening PCF and Sidney Linker and Josh Smith (Corning Optical Communications) for the data hall architecture design and product inputs.

## References

Herrmann, C., Diaz, A., Wu, T., & Von Zimmerman, R. (Mar 2026). *Product Carbon Footprint Screening of a Data Hall v3.1*. Sphera Solutions Inc. on behalf of Corning Optical Communications LLC.

Smith, W. (Jan 2026). *The Trade Shortage Crisis: How Data Centers Are Exposing America's Skills Gap*. Retrieved from Forbes: <https://www.forbes.com/councils/forbesbusinesscouncil/2026/01/23/the-trade-shortage-crisis-how-data-centers-are-exposing-americas-skills-gap/>

Sphera. (2026). *Sphera Life Cycle Assessment Datasets (MLC, Managed Life Cycle Content)*. Retrieved from <https://lcadatabase.sphera.com/>

United States Environmental Protection Agency. *Greenhouse Gas Equivalencies Calculator*. (2026). Retrieved from <https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator#results>