

Connecting the Future – from Cloud to Edge

Manja Thessin, RCDD, RTPM, Enterprise Market Manager, AFL

NOVEMBER 15, 2022

KEY TAKEAWAYS

- Rapidly evolving technology capabilities drive exponential growth in data.
- The future of enterprise computing is at the edge.
- Advancements in cable technology support connectivity needs of the future.
- AFL provides future-proof connectivity along the edge continuum.

in partnership with



Connecting the Future – from Cloud to Edge

OVERVIEW

As enterprise data centers move toward the edge, fiber infrastructure needs to be developed with the future of connectivity in mind. From virtual conferencing that drives point-of-presence deployments to autonomous vehicles that will potentially require compute power on city streets, the emerging edge will have a major impact on tomorrow's networks and data center infrastructure.

AFL, an end-to-end solutions provider of fiber optic cable, transmission and substation accessories, and more, produces a range of cabling, patching, and splicing solutions to meet the demands of next-generation applications at the edge and future-proof networks.

CONTEXT

Manja Thessin discussed trends in data center computing and connectivity, and explained the connectivity requirements for each area of the edge.

KEY TAKEAWAYS

Rapidly evolving technology capabilities drive exponential growth in data.

The pandemic accelerated digital adoption, as social changes such as remote work and learning fueled the need for high-speed internet connectivity, more online services, and improved streaming applications. The shift to a more connected society also sped innovations that leverage AI and ML and depend on advanced communication technologies, such as 5G, that increase bandwidth and decrease latency.

The real-time response and decision making enabled by high-speed broadband provide a foundation for new applications, such as augmented reality and autonomous vehicles, to be developed and adopted at

scale. However, these applications create and depend on exponentially growing volumes of data, in turn driving the need for a massive increase in storage and compute power.

The future of enterprise computing is at the edge.

Public and hybrid cloud solutions have been widely adopted in the enterprise data center. However, in cloud computing, the application environment is global, requiring significant network bandwidth along with high latency that cannot meet the sub-5-millisecond needs of some real-time applications. The volume of data that next-generation applications generate, coupled with the ultra-low latency requirements of AI- and ML-driven technology, is pulling processing capabilities ever closer to the edge.

Latency at the cloud data center measures at 100 milliseconds or more. Most of today's applications can tolerate this; however, such high latency is not sufficient for applications such as controlling a fleet of drones, piloting autonomous cars, or performing remote surgery. Edge compute is the physical compute infrastructure that is positioned on a spectrum between the hyperscale cloud and the device, or application edge, that supports a variety of applications.


To address latency issues, some businesses are building micro data centers and "cloudlets" (considered the "near-edge"), which are small, generally unstaffed facilities that allow highly localized placement of mission-critical infrastructure to bring latency down to around 20 milliseconds or faster.

Latency is the currency of the edge.

Manja Thessin, AFL

Connecting the Future – from Cloud to Edge

Figure 1: Latency decreases as computing moves closer to the edge



The emerging edge represents a dynamic change in the enterprise data center. By 2024, the enterprise edge is projected to grow about five times, from 50 million to 250 million device shipments. By shifting processing to the edge using highly localized robust compute architecture, the ultra-low latency demanded by applications such as autonomous vehicles and smart AI agriculture can be realized—without the demand for network bandwidth. This shift not only improves speed, but also lowers cost by reducing the volume of data transfer over the network.

Advancements in cable technology support connectivity needs of the future.


An unprecedented level of connectivity is needed to enable this new approach to computing, requiring investment in a deep, dense fabric of fiber and other infrastructure deployments to support resilience and reliability of increasing network demands. The cost and complexity of deploying networks that are truly future proof require solutions to be **expandable, flexible, and accessible**.

The ultimate goal of both hyperscale and edge data center providers is to maximize bandwidth per square inch. (Hyperscale does not refer to size, but rather ability to scale in response to increased demands.) That requires more than just fiber and cable innovation. Managing this level of density—where multiple cables, each with thousands of fibers, need to be patched and spliced—requires innovation in splice enclosure designs, patch panels, and fiber distribution management. Merging those innovations in cable design and high-density fiber management results in expandability, flexibility, and accessibility in a true end-to-end solution.

Size reduction is one of the greatest differentiators between solutions. Cable manufacturers know it is critical to reduce the size and weight of the cabling infrastructure. Today's flexible ribbon solutions can provide 864 fiber cables in a similar space to what was previously required for 144 fiber loose tube or 288 fiber traditional ribbon solutions. A 100-200% density increase can be achieved within the same space using new-design ribbon cables.

Connecting the Future – from Cloud to Edge

Figure 2: Next-generation applications pull computing power to the edge



Source: ATOM Micro DC Meets Edge (<https://atom.tech/wp-content/uploads/2019/01/Industry-Application-Guide-Micro-Data-Center-Meets-Edge-Computing-Atom-Technology.pdf>)


The new flexible ribbon design is a technical enabler of 200-micron single-mode fiber, which allows for a significant breakthrough in density per square inch. The 200-micron fiber provides over 35% reduction in overall fiber cross section, which translates to an ability to support higher density, less overall weight, and a smaller bend radius that makes for easier installation. Ribbon solutions that maintain a 250-micron pitch enable backward compatibility in existing data center facilities where the need to maximize connections is competing against the need to minimize the infrastructure size.

Beyond density, fiber manufacturers also place priority on improving flexibility in deployment. Modern gel-less ribbon cables combine the benefits of ribbon and loose tube, and produce no mess, saving installation time. Technicians can work with gel-less cables as a 12-fiber ribbon or as an individual fiber.

Along with the easier cable management resulting from smaller, lighter cables, and the ability to perform mass fusion splicing, these cabling advancements offer a potential 89% throughput improvement. And, modularity is important in high-density connectivity platforms to maximize bandwidth per square inch and drive continued economic benefit. Modules can be developed to support new and emerging technologies or architectures, allowing for different housings to be compatible with the same cassettes. This allows field teams to be trained on a common set of modules and cassettes, for providers and integrators to have fewer unit types to stock and manage in their inventory, and for manufacturers to achieve efficiency, supply chain resiliency, and opportunities for cost benefits.

Connecting the Future – from Cloud to Edge

Figure 3: Cable types and associated fiber count



AFL provides future-proof connectivity along the edge continuum.

AFL has defined key areas of connectivity in the modern IT infrastructure, from the large-scale data centers to the far edge. The “edge continuum” is composed of seven core areas, each with its own cabling and connectivity needs.

Our quest to digitize the world depends upon a vast infrastructure of fiber and optical connectors, coupled with a variety of panels and cabinets. When you choose the right components, a truly future-proof network can be achieved that can support the transformative applications of our emerging reality.

Manja Thessin, AFL

At the heart of data center infrastructure sits the public cloud (E6), dominated by hyperscalars such as Google, Amazon, and Microsoft. These massive facilities are typically located in sparsely populated areas, holding over 5,000 servers in an area of over 10,000 square feet, with millions of connections.

The regional edge (E5) is where most edge deployments are happening today, where content distribution network providers are looking to distribute their locations and meet increased demand for data centers at a more local level. These are predominantly colocation data centers that host hundreds of different customers and can achieve economies of scale, sharing cooling power, operations, and maintenance expenses. Customers rent cabinets, server racks, cages filled with cabinets, or even the entire data center. Telco service providers locate their core interfacing nodes and provide customer connections to their access network via the regional edge. On average, latency is reduced by half over the cloud data center. E5 is where the most investment is made in edge data centers today, with typical infrastructure including ultra-high fiber count foldable ribbon cables and high-density splice and patch frames.

The network edge (E4) is where edge nodes reside, mostly in central offices. These are data center-like environments in which companies leverage existing facilities in their network for edge compute possibilities. Typically in every locality, these are privately owned by telco service providers and considered the outer edge of the mobile network, and can also be referred to as the “carrier edge” of the network—traditionally where telephone lines in a region would meet and be redirected to a destination.

Connecting the Future – from Cloud to Edge


E4 requires a solution that blends existing infrastructure with the compute power necessary to support emerging applications, requiring optical solutions such as modern flexible ribbon with a 250-micron pitch, important for matching the fiber pitch of existing central office optical infrastructure. Suitable options include bulk cable terminated with field-installed MPOs and dedicated patching and splicing closures or pre-terminated cabling. Prefabricated edge racks are increasingly in use in E4.

At some point, the edge network will need to extend to offices, campuses, factories, venues such as stadiums, hospitals, and logistics centers to support data collection from the millions of IoT devices and sensors. These special-purpose edge data centers, or enterprise edge (E3), exist in different locations and configuration, from traditional racks within an existing on-premise data

center, to small-scale ruggedized standalone solutions that can operate in harsh environments, and a wide range in between—resulting in a range of infrastructure requirements, including armored cables, wall-mounted splice enclosures, and data center vaults.

E2 is where edge compute begins to blend with the wireless core to support applications such as smart agriculture, where cellular base stations serve as key points to connect end-user devices to the core network. E2 edge compute is co-located with small cells, centralized radio access network (CRAN), fixed wireless access (FWA), or a private LTE cell node. These are challenging environments that require the capability to meet strict space constraints, as well as to operate in dusty, less maintained and lower-temperature-regulated conditions.

Figure 5: The edge continuum



Connecting the Future – from Cloud to Edge

The most advanced applications enabled by 5G will depend on infrastructure such as street-side cabinets, which could become a key data center form factor at the edge, holding anywhere from a quarter rack to two full racks of modern data center equipment in highly remote locations. As these future solutions are developed to support applications requiring even greater density, it is possible that compute power will be integrated into the optical infrastructure itself, such as aerial and underground splice enclosures (AFT patent pending).

These are the types of innovations that we need to begin to be thinking about today, even though the applications that will require them are not here yet, because if we don't start today, we will likely not get there.

Manja Thessin, AFL

The E6 and E5 spaces are driving the ultra-high fiber count cable and connectivity demand, with demand growth for fiber kilometer estimated at well over 13% over the next five years, while carrier deployments in support of 5G and FTX provide the backbone for the edge. As those infrastructures are deployed, it is critical to consider and plan for expandability, flexibility, and accessibility that will allow for edge compute capability in other areas to be added later.

ADDITIONAL INFORMATION

- AFL. To learn more, visit www.aflglobal.com

BIOGRAPHY

Manja Thessin, RCDD, RTPM

Enterprise Market Manager, AFL

Manja Thessin serves as Enterprise Market Manager for AFL, leading strategic planning, and market analysis initiatives. Manja has 22 years of ICT experience in the field, design & engineering, and project management. She has managed complex initiatives in Data Center, Education, Industrial/Manufacturing, and Healthcare. Manja earned a master's certificate from Michigan State University in Strategic Leadership and holds RCDD and RTPM certifications with BICSI®.