WHAT’S NEXT FOR THE DATA CENTER

2024 trends to watch

Contents

About the authors

Dr. Earl Parsons

Dr. Earl Parsons is the Director of Data Center Architecture Evolution at CommScope. He joined CommScope in 2014 as a Principal Optical Engineer.

Prior to joining CommScope Earl received an MS and Ph.D. in optical sciences from the University of Arizona and was a Senior Member of Technical Staff at TE SubCom, formerly known as Tyco Telecommunications. Dr. Parsons also served as an editor of the IEEE 802.3db-2022 standard. His interests include multimode and single mode fiber optic systems to enable artificial intelligence data centers.

Kevin Swank

Kevin Swank is a veteran of the cellular wireless industry – having worked in everything from 2G to 5G. He started his career as a software engineer and moved into leadership positions in product management and product marketing. Having worked for NEC, Nortel Networks, InnerWireless, and Black Box, Kevin is currently responsible for product marketing direction for Intelligent Cellular Networks at CommScope.

Ken Hall

Ken Hall is data center architect for North America at CommScope, responsible for technology and thought leadership as well as optical infrastructure planning for global scale and related data centers. In this role he has been instrumental in the development and release of high speed, ultra low loss fiber optic solutions to efficiently enable network migration for data center operators.

Previously Ken worked with TE Connectivity/ Tyco Electronics/AMP in a variety of roles. His experience includes global Network OEM and Data Center program management and strategy, project management, marketing, industry standards and technical sales management. Ken was also responsible for industry standardization and proliferation of copper and fiber small form factor connectors and high-density interfaces for network electronics OEMs.

Ken has nine patents to date for fiber-optic connectors and infrastructure management systems.

Ken graduated with a Bachelor of Science from Shippensburg University. He is a registered Communication Distribution Designer (RCDD) and Network Technology Systems Designer(NTS).

Alastair Waite

Alastair Waite joined CommScope in September 2003 as a Product Manager for the company's Enterprise Fibre Optic division, since that time he has held a number of key roles in the business including Head of Enterprise Product Management, for EMEA, Head of Market Management and the Data Center business leader in EMEA.

Since January 2016, Alastair has had responsibility for architecting CommScope's data center solutions, ensuring that customer's infrastructures are positioned to grow as their operational needs expand in this dynamic segment of the market.

Prior to joining CommScope, Alastair was a Senior Product Line Manager for Optical Silicon at Conexant Semiconductor, where he had global responsibility for all of the company's optical interface products.

Alastair has a BSc in Electronic Engineering from UC Wales.

It's a brave new world. Here's what you need to know to roll with the changes.

Bob Dylan said it best: Oh, the times, they are a-changin'. Truer words have never been spoken, especially if you're a data center manager. Overnight, or so it seems, application demands went from a cautious yellow alert to flashing bright red.

2023 saw the unleashing of artificial intelligence in all its guises—generative AI, machine learning, artificial neural networks, deep learning, natural language processing and more. In its wake, data center managers and their teams are scrambling to figure out how to handle not just the additional petabytes of new data flooding their networks, but the ultra-low latency requirements, increases in power usage, and exponentially higher fiber counts that are needed.

And the application demands aren't the only things increasing. Data center footprints are growing like kudzu, extending beyond national borders. Managing data sovereignty has become as much a political challenge as it is a legal and security concern.

At the same time, data centers of all kinds—hyperscale, cloud providers, MTDCs and enterprise—are forming new relationships to take advantage of prime market locations, available resources and much-needed white space. As a result, the task of figuring out the DCI backbone cabling has become a full-time job itself.

Of course, this is all taking place against the backdrop of a global climate crisis that has data center operators rethinking how each step in their supply chain and network lifecycle affects their organization's sustainability targets.

The times are a-changin' all right, and the common denominator is that all these changes have a profound impact on your network's infrastructure—from its cabling, connectivity and components to its architecture, resilience and adaptability. In the chapters that follow, we explain the issues you face and offer innovative solutions that showcase CommScope's deep experience and out-of-the-box approach.

Welcome to CommScope's most recent edition of the Data center trends e-book. Embrace the change!

1

Adapting to higher fiber counts in the data center

The ongoing challenge for any data center operator is how to provide ever-increasing capacity. The ability to support higher data rates, port counts and fiber density is key to increasing capacity and, ultimately, the customer's experience. That's the goal, right? But there is no magic bullet, no single solution that can completely address the capacity issue. However, by combining related technologies in a planful way, data center operators can support continuous capacity growth and enable best-in-class customer experience. In this chapter, we'll discuss the issues and trends driving the search for repeatable capacity growth.

Congested cable pathways

To achieve predictable and scalable capacity growth, the network infrastructure building blocks—including cable, pathways, connectivity and network architectures—must be aligned to deliver that capacity efficiently and provide the most value. As applications have evolved, counts for singlemode and multimode fiber have progressed from 2 to 8 and 16 fibers per switch port. While fiber counts have increased by at least 4x or 8x, the pathways and spaces between and around the equipment have not. Therefore, any capacity solution must take into account innovations in fiber cable construction and connectivity that make more efficient use of the available space.

Compact cable construction—when combined with higher fiber counts—is especially useful when interconnecting data centers. Data center interconnect (DCI) trunk cabling with 3,000+ fibers is common for connecting two hyperscale facilities, and some operators are already doubling that designed capacity. Inside the data center, problem areas include backbone trunk cables that run between high-end core switches or from meet-me rooms to cabinet-row spine switches.

Rethinking the fiber package

The progression of fiber and optical network has been a continual response to the need for faster, bigger data pipes. As those needs intensify, the ways in which fiber is designed and packaged within the cable have evolved—allowing data centers to increase the number of fibers in a cable construction without necessarily increasing the cabling footprint. Rollable ribbon fiber cabling is one of the more recent links in this chain of innovation.

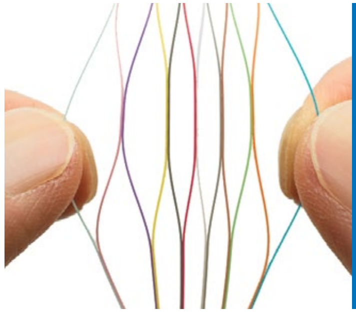

Rollable ribbon 250 μm

- Partially bonded individual 250-micron fibers
- Reduces time using mass fusion splicing
- 20-40 percent smaller cable OD over matrix ribbon, offering better duct utilization

Rollable ribbon fiber cable is based, in part, on the earlier development of the central tube ribbon cable. Introduced in the mid-1990s, primarily for outside plant (OSP) networks, the central tube ribbon cable featured ribbon stacks of up to 864 fibers within a single, central buffer tube. The fibers are grouped and continuously bonded down the length of the cable, which increases its rigidity. While this has little effect when deploying the cable in an OSP application, in a data center a rigid cable limits the more flexible routing which is required to navigate narrow and congested pathways.

In a rollable ribbon fiber cable, the optical fibers are attached intermittently to form a loose web. This configuration makes the ribbon more flexible, allowing as many as 3,456 fibers to be loaded into one two-inch duct—twice the density of conventionally packed fibers. This construction reduces the bend radius, making these cables easier to work with inside the tighter confines of the data center.

Inside the cable, the intermittently bonded fibers take on the physical characteristics of loose fibers that easily flex and bend—making it easier to manage in tight spaces. In addition, rollable ribbon fiber cabling uses a completely gel-free design, which helps reduce the time required to prepare for splicing, therefore reducing labor costs. The intermittent bonding maintains the fiber alignment required for typical mass fusion ribbon splicing.

Reducing cable diameters

For decades, nearly all telecom optical fiber has had a nominal coating diameter of 250 microns. With growing demand for smaller cables, that has started to change. Many cable designs have reached practical limits for diameter reduction with standard fiber. But a smaller fiber allows additional reductions. Fibers with 200-micron coatings are now being used in rollable ribbon fiber and micro-duct cable.

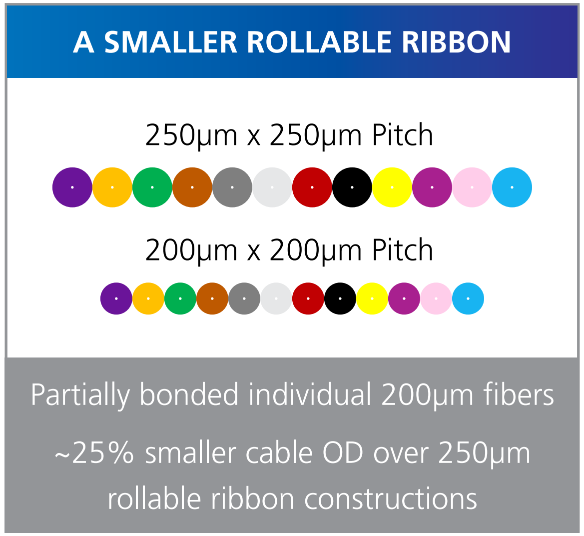

It is important to emphasize that the buffer coating is the only part of the fiber that has been altered. 200-micron fibers retain the 125-micron core/cladding diameter of conventional fibers for compatibility in splicing operations. Once the buffer coating has been stripped, the splice procedure for 200-micron fiber is the same as for its 250-micron counterpart. Fixtures are also available to enable "up-pitching" from 200-micron to 250-micron for the transition to multifiber connectors, if necessary, at the endpoints.

While duplex applications can be supported by any of the available 8-, 12-, 16- or 24-fiber subunit counts, migration to 8- or 16-fiber applications is best supported by MPO8 or MPO16 fiber trunk units. Should 16-fiber applications be the network team's current or possible future plan, 16-fiber trunks will provide the most efficient Day 1 installation. This configuration can support all existing applications without wasting fibers or needing to bridge trunk cables on site in the future. Additional benefits from these reduced-diameter cable constructions include the associated reduction of pathway space required for the fiber count and reduced cabling materials used—delivering sustainability values.

New smaller connectors support higher fiber counts, easier installation

Connectors are also evolving, with new VSFF (very small form factor) configurations available to provide duplex or parallel application support. Common fiber counts for higher speed connectors include 2, 8 or 16 fibers in both singlemode and multimode. For some applications, 24-fiber cabling remains an option as well. Decisions on which cable subunits to use for the trunks should consider current and possible future connector requirements.

Recently-introduced VSFF duplex connectors aligned with transceivers have entered the market over the past several years. They provide better density and, in some cases, breakout options directly at the transceiver. The intent is to enable higher fiber counts to enable full capacity utilization at the network equipment. Shown below for size reference are the legacy LC duplex along with SN, MDC and CS connectors. While aligning with transceiver applications, they also can provide manageable higher density at the patch panel for structured cabling applications.

Other soon-to-be-available multifiber options in the SN and MDC footprint will pack even more manageable fibers in a smaller size. These Very Small Form Factor (VSFF) connectors house 16 or 24 fibers in the same space as the SN and MDC duplex solutions.

MMC16 & MMC24 – USConec Ltd

SN-MT16 & SN-MT24
SENKO Advanced Components

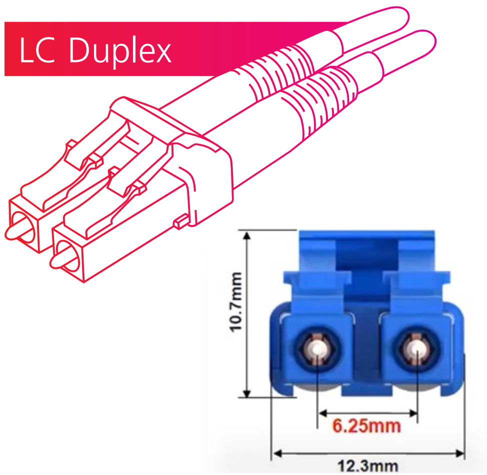

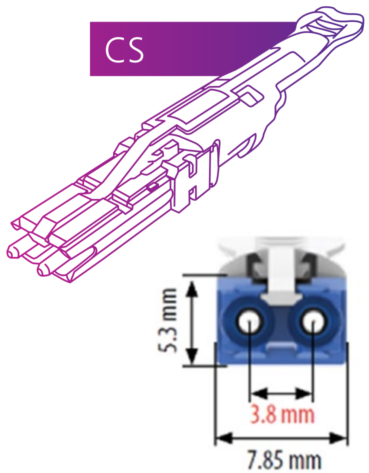

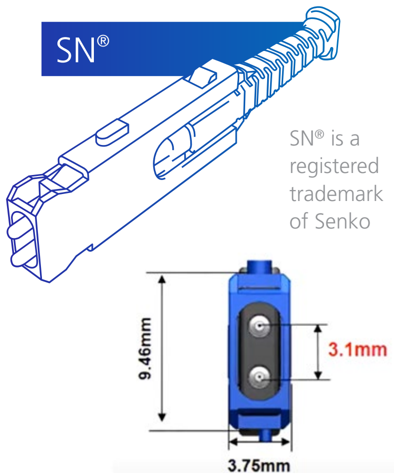

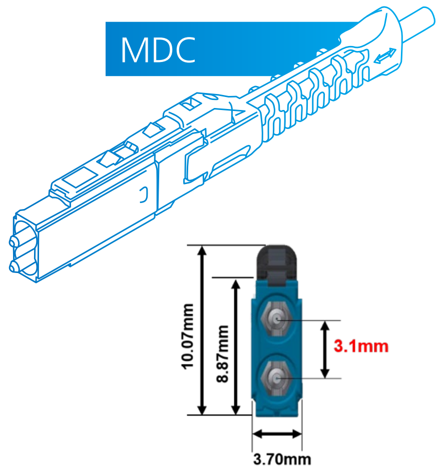

The enabling technology is a smaller type of MT ferrule (similar to those found in MPO connectors) that can house more fibers within the highly compact connector body. These connectors are not intermateable, but are currently developing in the market.

Additionally, VSFF fiber connectors enable the smaller and lighter pre-terminated high fiber-count trunks to be pulled through conduits more easily. When deployed using low-profile rollable ribbon cable, pre-terminated high fiber-count VSFF connectors simplify and accelerate installation, saving valuable time and space; having been assembled in a carefully controlled factory setting it also provides added performance assurance. As a result, rollable ribbon cable pre-terminated with the new VSFF connectors provide some unique benefits compared to field terminated or spliced alternatives.

Not surprisingly, initial applications for VSFF connectors target pre-terminated high fiber-count trunk cables designed to be pulled through innerducts or raceway systems. Future VSFF applications could include equipment interface, breakouts or structured cabling.

For channel design and application cabling considerations, please contact your local CommScope sales engineer or reference the Propel™ Design Guide.

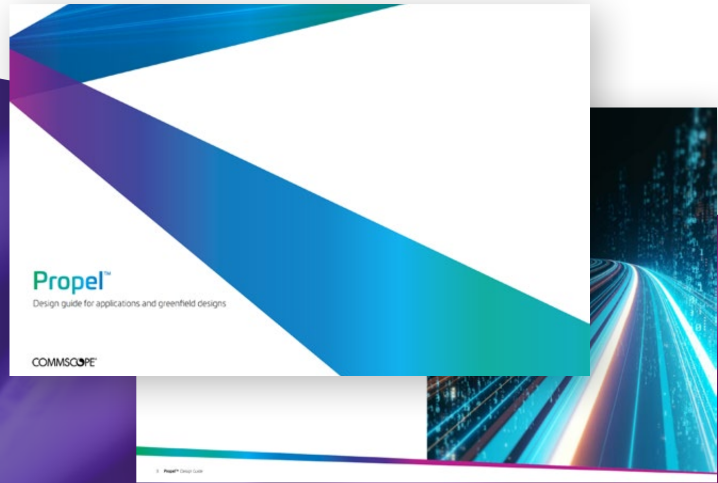

2

400G, 800G, and 1.6T data center links

Introduction

The first measure of an organization's success is its ability to adapt to changes in its environment. Call it survivability. If you can't make the leap to the new status quo, your customers will leave you behind.

For cloud-scale data centers, the ability to adapt and survive is tested every year as increasing demands for bandwidth, capacity, and lower latency fuel migration to faster network speeds. During the past several years, we've seen network fabric link speeds throughout the data center increase from 25G/100G to 100G/400G. Every leap to a higher speed is followed by a brief plateau before data center managers need to prepare for the next jump.

Currently, cloud and hyperscale data centers are deploying links with 800G transceivers, while the industry seeks to standardize both 800G and 1.6T transceivers. A key consideration is which optical technology is best. Here, we break down some of the considerations, tradeoffs and options of 400G, 800G and 1.6T optical transceivers.

Optical transceiver types

Optical transceivers can be grouped by their supported reach and fiber type. SR optics typically support 100 m reaches over parallel multimode fiber. DR optics use parallel singlemode fiber up to either 500 m or 2 km in length. And FR and LR optics use duplex singlemode fiber and wavelength division multiplexing (WDM) up to 2 km and 10 km, respectively. These names are consistently used for 400G, 800G, and 1.6T transceivers.

400G optical transceivers

The optical market for 400G is being driven by cost and performance, and transceivers using 4x100G lanes are replacing the earlier iterations that used 8x50G lanes. 4x100G switches and transceivers offer lower cost and power consumption compared to 8x50G transceivers.

There are several options for optics using parallel fiber at 400G. The IEEE 802.3db standard codified 400G transceivers using parallel multimode fiber. Optics compliant to 400G-SR4 will support 100 m over eight fibers (OM4 or OM5). A new application dubbed 400G-VR4 will support 50 m reach of eight OM4 or OM5 fibers and targets in-row applications. For parallel singlemode fiber, 400G-DR4 and 400G-DR4-2 will use 8 fibers up to 500 m or 2 km in length, respectively.

Each of the parallel fiber options in Table 1 use eight fibers, but because the transceivers accept either an MPO8 or MPO12 connector, only the outer eight fibers are used. This follows a multigenerational trend in which the middle four fibers in a 12-fiber cable are not used. This has prompted cable companies like CommScope to introduce a line of eight fiber cables that only include the eight fibers used for transmission. When terminated with MPO8 connectors, the eight-fiber cables can be intermated with MPO12 cables and are compatible with all 400G parallel optics.

Table 1

400G parallel optics

Application	Reach	Fiber #	Fiber type
400G-SR4	100 m	8	OM4, OM5
400G-VR4	50 m	8	OM4, OM5
400G-DR4	500 m	8	SMF
400G-DR4-2	2 km	8	SMF

As shown in Table 2, 400G-FR4 and 400G-LR4 were standardized in IEEE 802.3cu and use WDM on duplex fiber. These optics combine four wavelengths on a single transmit-and-receive fiber. The FR optics support 2 km reach and offer an upgrade path from 100G-CWDM4. The LR optics reach up to 6 km if only compliant to 802.3cu (labeled 400G-LR4-6) and 10 km if compliant to the 400G-LR4-10 MSA specifications.

Table 2

400G WDM optics

Application	Reach	Fiber #	Fiber type
400G-FR4	2 km	2	SMF
400G-LR4	10 km	2	SMF

Beginning with the 400G generation, demand for optics with parallel fiber has grown faster than those using WDM. Optics for parallel fiber tend to offer lower cost and power consumption than WDM optics. Parallel fiber also offers data center operators more flexibility by enabling fiber breakouts, for example. The technology also allows a 400G transceiver with eight fibers on one end to connect to four different 100G transceivers using duplex fiber. Parallel fiber optics are being used for longer links, but WDM optics are not being used for shorter links.

800G optical transceivers

The first generation of 800G transceivers will use 8x100G lanes and parallel fiber. These transceivers will build on 400G transceiver technology and will be included in the IEEE 802.3df standard scheduled to be published in 2024. The IEEE 802.3dj standard project, scheduled to be published in 2026, will address optics running on 4x200G lanes.

Table 3

800G parallel optics

Application	Reach	Fiber #	Fiber type
800G-VR8	50 m	16	OM4/OM5
800G-SR8	100 m	16	OM4/OM5
800G-DR8	500 m	16	SMF
800G-DR4	500 m	8	SMF
800G-DR8-2	2 km	16	SMF
800G-DR4-2	2 km	8	SMF

Table 3 lists parallel fiber optics including 800G-VR8, 800G-SR8, 800G-DR8 and 800G-DR8-2. They are designed to support 50 m OM4 or OM5, 100 m OM4 or OM5, 500 m singlemode fiber, and 2 km singlemode fiber reaches, respectively. Each of these optics will require eight fiber pairs (16 fibers total) for 800G transmission. The connector interface will consist of MPO16 or two MPO8 connectors. The earliest adopters of 800G transceivers will use them as 2x400G, with each 800G transceiver behaving as two distinct 400G transceivers. It makes sense in these cases to have two MPO8 connectors at the interface to support this distinction. In the future, native 800G transceivers will be more dominant and will use the MPO16 connector. Nearly all transceiver manufacturers will offer both MPO16 and 2xMPO8 800G transceivers.

Table 4

800G BiDi optics

Application	Reach	Fiber #	Fiber type
800G-VR4.2	50 m/70 m	8	OM4/OM5
800G-SR4.2	70 m/100 m	8	OM4/OM5

The Terabit BiDi MSA released a specification for 800G transceivers using eight multimode fibers (see Table 4). These transceivers (800G-VR4.2 and 800G-SR4.2) will support either 50 m or 70 m over OM4 and 70 m and 100 m with OM5. These optics operate bidirectionally, with each fiber transmitting and receiving simultaneously. VCSELs of different wavelengths are used to generate separable transmit and receive signals. OM5 is the only multimode fiber specified to support multiwavelength operation, which is why it offers longer reach. BiDi enables 800G transceivers with 100G lanes to operate using only eight fibers.

Among the duplex singlemode options, there are also 800G optics that are really 2x400G-FR4. These transceivers require four fibers at the connector interface and will typically use two duplex LC connectors belly-to-belly. These transceivers are also good candidates for very small form factor (VSFF) connectors like SN or MDC.

Once IEEE 802.3dj standardizes 200G lanes, singlemode transceivers—including 800G-DR4, 800G-DR4-2, 800G-FR4 and 800G-LR4—will be specified. Like their 400G counterparts, the DR optics will use eight fibers up to 500 m or 2 km; the FR will use duplex fiber up to 2 km; and LR will use duplex fiber up to 10 km (see Table 5). These optics will be a drop-in replacement for 400G and use the same cable plant as 400G.

Table 5

800G WDM optics

Application	Reach	Fiber #	Fiber type
800G-FR4	2 km	2	SMF
800G-LR4	10 km	2	SMF

To date, IEEE 802.3 has not set objectives for 200G VCSELs and multimode fiber. This is not surprising and is consistent with previous generations. For each new speed, singlemode transceivers are specified first as it is easier to achieve high speed signaling with these more complicated transceivers. Multimode optics are lower cost and lower power, and their standards typically take more time to develop. We are confident that 200G VCSELs will be standardized in a future project.

1.6T optical transceivers

Beyond 800G, the next-generation transceiver will use the letter “T” for terabit per second. Doubling the data rate from 800 gigabits per second leads to 1600 gigabits (or 1.6 terabits) per second.

Table 6

1.6T BiDi optics

Application	Reach	Fiber #	Fiber type
1.6T-VR4.2	50 m/70 m	16	OM4/OM5
1.6T-SR4.2	70 m/100 m	16	OM4/OM5

The Terabit BiDi MSA has already specified 1.6T transceivers that use multiwavelength VCSELs and multimode fiber. As shown in Table 6, 1.6T-VR8.2 will support 50 m over OM4 and 70 m over OM5. For longer reaches, 1.6T-SR8.2 will support 70 m with OM4 and 100 m with OM5. Both these transceiver types will use 100G lanes and 16 bidirectional fibers. The MSA calls out MPO16 connector interfaces as the preferred connector.

Table 7

1.6T parallel optics

Application	Reach	Fiber #	Fiber type
1.6T-DR8	500 m	16	SMF
1.6T-DR8-2	2 km	16	SMF

Early singlemode 1.6T optics will be 1.6T-DR8 and 1.6T-DR8-2. Both will use 16 parallel singlemode fibers (or eight fiber pairs) and will support 500 m and 2 km reaches, respectively (see Table 7). These optics will use the same 200G lanes developed for 800G but will increase the lane count to eight. Like 800G-DR8, these 1.6T transceivers will likely be available with MPO16 and 2xMPO8 connector interfaces.

Multimode vs. singlemode

For links less than 100 m, data center operators have a choice: Do they deploy singlemode or multimode optics? While some in the industry consider singlemode fiber to be the more future-proof option, multimode offers many advantages. Data centers that deploy the latest speed transceivers will pay twice as much for a singlemode optic than an equivalent multimode optic. Over time, the price difference will converge; multimode transceivers are only slightly less expensive than singlemode when the speed is a few generations old. One thing that will not change is the amount of power consumed. Multimode transceivers consume 1-2 W less energy than singlemode. That leads to 2-4 W power savings per link.

Over time, transceiver technology will migrate to shorter and shorter reaches. As lane speeds increase, we see multimode fiber replacing copper cables for in-row applications. An obvious example is the use of multimode fiber in artificial intelligence (AI) clusters. We anticipate multimode fiber to remain a key part of data center networks for many more generations.

Conclusions

Transceiver innovation continues at a fast pace. The transceivers selected for a data center impact the fiber cable and connectors needed. CommScope is actively engaged with the transceiver ecosystem to ensure that our customers have the right connectivity for their networks. Read more about steps you can take today to ensure your fiber infrastructure is ready for this future at commscope.com.

3

Addressing the power challenge

Data center operators often take different approaches to network design and platforms. However, every operator, regardless of size, has one thing in common: power—how to get it and how to conserve it. Now, for the first time in the history of the data center industry, the ability to deliver power to the right place and at the right price can no longer be guaranteed. In fact, the decision of where to build the data center is governed as much by these external factors as it is by the operator's business strategy.

In this chapter, we'll explore three primary factors that can challenge an organization's plans when it comes to power.

Power supply issues

In a 2022 earnings conference call, Andy Power, president and CEO of Digital Realty—the world's largest MTDC—revealed that a primary electricity provider in Northern Virginia had warned customers of a bottleneck in eastern Loudon County that could delay deliveries until 2026[i]. The issue was not reported as a power generation problem, but rather a lack of transmission lines to carry the power required.

Sometimes, the forecasted demand for data center power is enough to prompt governmental action. Following alerts from energy providers, the Office of the Mayor of London (UK) published a briefing paper describing a rapid influx of requests for new electricity connections throughout West London. Most of the new requests were from data center operators seeking to colocate adjacent to fiber-optic cables that pass through the region along the M4 corridor.

According to the Mayor's office:

"Data centers use large quantities of electricity, the equivalent of towns or small cities, to power servers and ensure resilience in service. The scale of electricity demanded by these data centers has created capacity constraints on both the distribution and transmission networks in the region, absorbing remaining electricity capacity in…(the) West London region for the remainder of the decade…major new applicants to the distribution network, including housing developments, commercial premises and industrial activities, will have to wait several years to receive new electricity connections."

– Mayor of London Office, GLA[ii]

Politics

Recently, the country of Ireland has become a hotbed of data center building activity. Their success is something of a double-edged sword. The Irish government must balance its leading position in the global data center industry with its responsibility to supply energy to its citizens and protect the environment. Increasingly, the Irish parliament is looking at what can be done and what part it should play in regulating data center builds.

From 2015–2022, the power consumed by data centers in the Republic of Ireland (5,200 GWh) increased 400 percent and, in 2022, represented 20 percent of all power generated in the country.

Central Statistics Office of Ireland, 2023

Simon Coveney, Irish Minister for Enterprise, Trade and Employment, said: "We clearly have some challenges over the next two to three years in terms of the energy demand…you have to manage data, and that involves data storage and data centers. The challenge is to find a way of powering them with sustainable and abundant power by capturing the potential, in particular, of offshore wind."[iii]

Moreover, Social Democrats TD and spokeswoman for climate, Jennifer Whitmore, suggests the possibility of forcing data centers to operate more efficiently: "What needs to happen is a moratorium on the granting and connection of data centers until there has been a strategic review around data centers. This would need to look at how they can be more efficient. There is no real oversight on this, which is a difficulty."

The following chart illustrates the trajectory for data center power consumption globally, and in particular server power consumption—clearly supporting the need for highly efficient data center designs as a key instrument to keeping federal regulators onside.

Annual energy consumption per end-user category (TWh), 2016-30

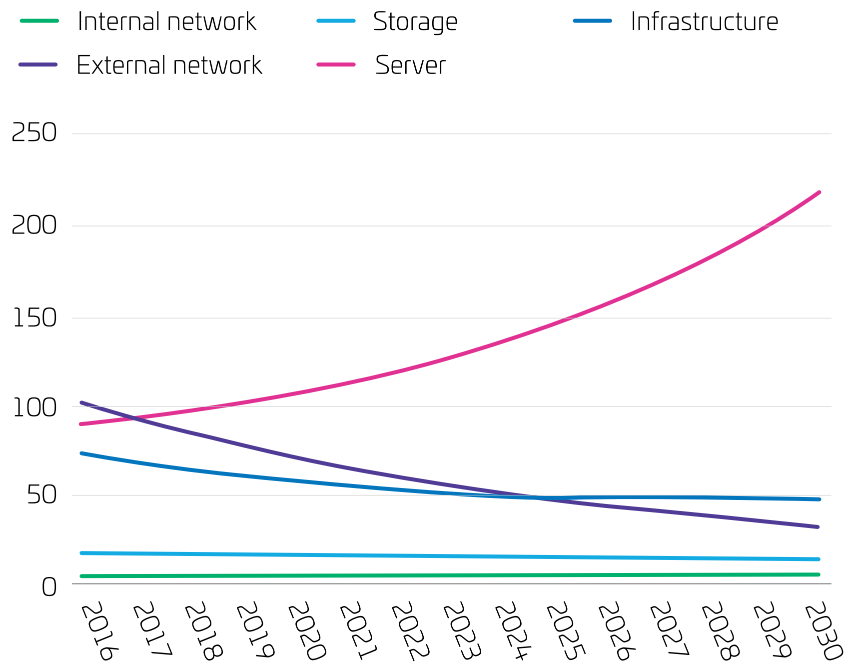

Figure 1
Source: Techmonitor

In Eastern Europe, the political decisions that have led to the current conflict in Ukraine have had a huge impact on the data center industry in the region. Many Eastern European countries had previously benefitted from their geographic locations at the confluence of the West, Russia and Asia. Those governments are now having to revisit their raison d'etre in the digital economy.

At the same time, the Russian/Ukrainian conflict has significantly disrupted the distribution and pricing of the natural gas and oil upon which much of Eastern and Western Europe have relied on. Prior to Russia's invasion of Ukraine, much of Europe's demand for gas and oil was filled by Russia. Since the invasion of Ukraine, natural gas prices have increased dramatically. Before the war, countries like Germany—which was sourcing half of its natural gas and around one-third of its oil from Russia and had some of the highest energy prices in Europe—saw energy prices rise even further. Sanctions applied by both sides have resulted in Germany looking to other methods of generation and energy supply, leading to further cost increases for homes and businesses alike. Figure 2 shows the sudden and dramatic increase in the commercial prices of natural gas across the European Union.

It is interesting to note that the upward trend in the EU's natural gas prices pre-dates the onset of Russian's aggression against Ukraine by about a year. Since the last quarter of 2022, energy prices have stabilized due to a combination of government intervention and supply chain efficiencies.

Development of natural gas prices for non-household consumers, EU, 2008-2022

(€ per kWh)

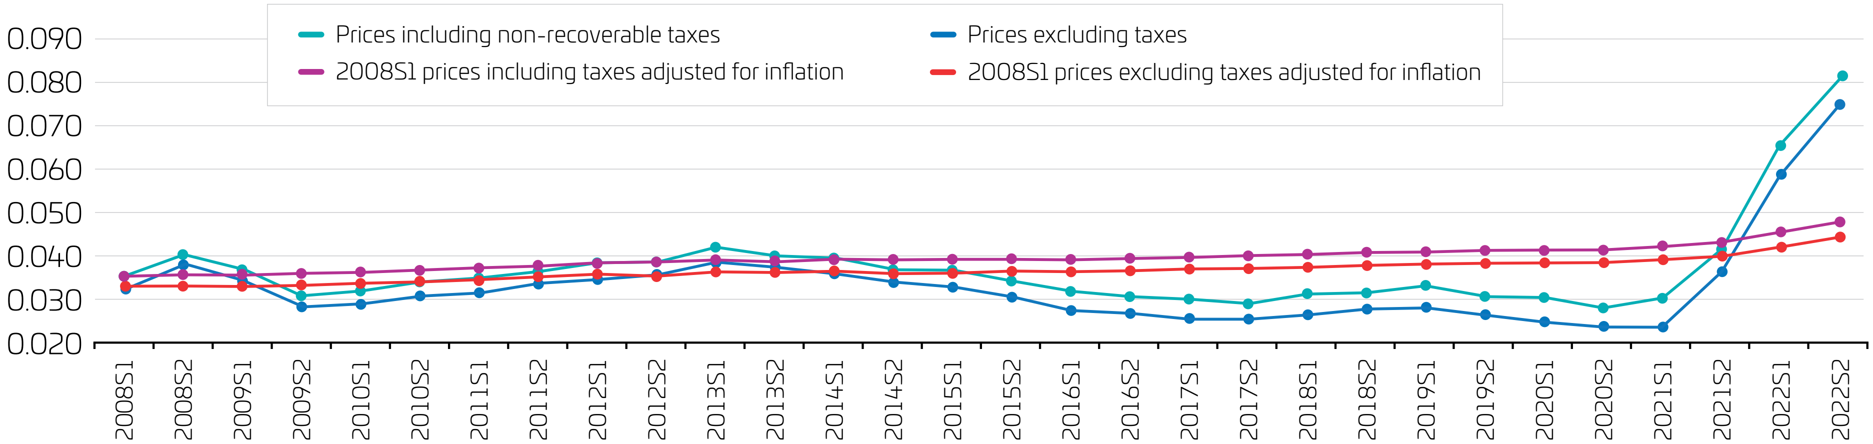

Figure 2 Source: Eurostat (nrg_pc_203)

Social demography

The third and final external lever affecting power availability and affordability for data centers is a shift in the world's population centers. With regard to choosing optimal data center locations, conventional business thinking says go after the low-hanging fruit first. For many years, that meant investing in a select number of markets in the United States and Western Europe. The U.S.—with its large middle class, high disposable income, and relatively few high-density metro areas—represented the easiest pickings. It also offered data center operators the conveniences of a common language, regulatory system, supply chain routes and a history of delivering large-scale CapEx builds.

The Europe market was a bit more difficult, but hardly impregnable. The idea was to start with Tier 1 cities, where power is readily available. The challenge was in navigating the different national governments, local business practices and languages. But given that Western Europe is home to a large middle class and a multitude of well-financed, international corporations, the reward was worth the effort. However, the cost of land and power in these European cities commands a high premium from data center builds, as organizations are having to settle for Tier 2 locations like Madrid, Marseille, Milan, Zurich, Berlin and Stockholm. Instead of focusing on a few key high-value Tier 1 cities, data center operators must widen the search, which only adds more intra-continental complexity. This means managing operations across four additional countries and power grids, and three more languages.

While it is easy to intellectualize the complexities of adding a large data center to an operator's existing portfolio, it can be difficult to truly grasp the challenges awaiting hyperscale and cloud providers that are looking to grow their footprint.

Tables 1 and 2 help put the issues in context. Table 1 shows the population and MTDC power available in each of the world's seven strongest economies: the G7. In these prime markets, the average MTDC power demand per million people is a whopping 48.7 megawatts (MW). This is the IT power load required to support a world-class digital economy.

Table 2 compares the same metrics for some of the fastest growing markets in Asia, which also happens to be the fastest growing MTDC region. Even with Singapore and Australia positively skewing the metrics for the region, the average MW IT load per head of population among those seven countries is less than one-tenth of that of the G7 countries.

Based on these and other findings from across the APAC region, it is clear that localized growth in data, population and 5G delivery demand is outpacing growth of data center capacity. This suggests that aggressive build initiatives—at both the government and private level—are needed if data center capacity is to keep pace with the population and demand for more data.

From a population and cost standpoint, the APAC region offers an attractive upside for new data center construction. It is also highly challenging. As difficult as it is to move a new-build project from the U.S. to Tier 1 and Tier 2 cities in Europe, consider how much tougher it is to try and deploy that same capacity in a country where power generation and supply are underdeveloped and there might not be enough water supply to support a rapidly expanding local population. Efficiency is key.

G7 Country	Population (Millions)	Data Center Market Size	Data Centre MW's/Population
USA	335	26,000	77.6
Japan	125	3,600	28.8
Germany	84	2,000	23.8
France	68	1,500	22.1
UK	67	3,000	44.8
Italy	59	500	8.5
Canada	40	1,300	32.5
Total	**778**	**37,900**	**48.7**

Table 1

APAC Country	Population (Millions)	Data Center Market Size	Data Centre MW's/Population
India	1,392	1,700	1
Indonesia	278	600	2
Philippines	111	300	3
Vietnam	100	315	3
Malaysia	33	600	18
Australia	26	1,100	42
Singapore	5	1,000	200
Total	**1,945**	**5,615**	**3**

Table 2

Source: 451 Research & CommScope

The role of the MTDC in addressing the power challenge

Given the costs and uncertainties inherent in local power supplies, political instabilities and shifting demographics, data center operators are rethinking their vertical integration strategies and the wisdom of building their own facilities. This is particularly true of organizations in which the data center is a support function as opposed to the primary revenue generator (think soft-drink manufacturers versus cloud services providers). As a result, more and more cloud-based and hyperscale operators are opting instead to partner with MTDCs that have existing capacity around the world.

In many ways, this new partner model offers more benefits than challenges. MTDC operators are real estate savvy and are optimized to satisfy tenants' evolving demands for world-class white space and reliable and affordable power. Perhaps more importantly, MTDC facilities are located in prime metro areas—perfect for cloud-based and hyperscale operators that need to support low latency and ultra-low latency mobile edge compute instances for 5G, industry 4.0 and IoT applications. Best of all, these facilities already exist, enabling larger data centers to roll out services quickly and easily with a high and faster return on investment. Figure 3 illustrates how the servers, and therefore the whole operation, can be spread across a multitude of locations in an effort to provide edge computing. MTDCs play a part in supporting this rollout.

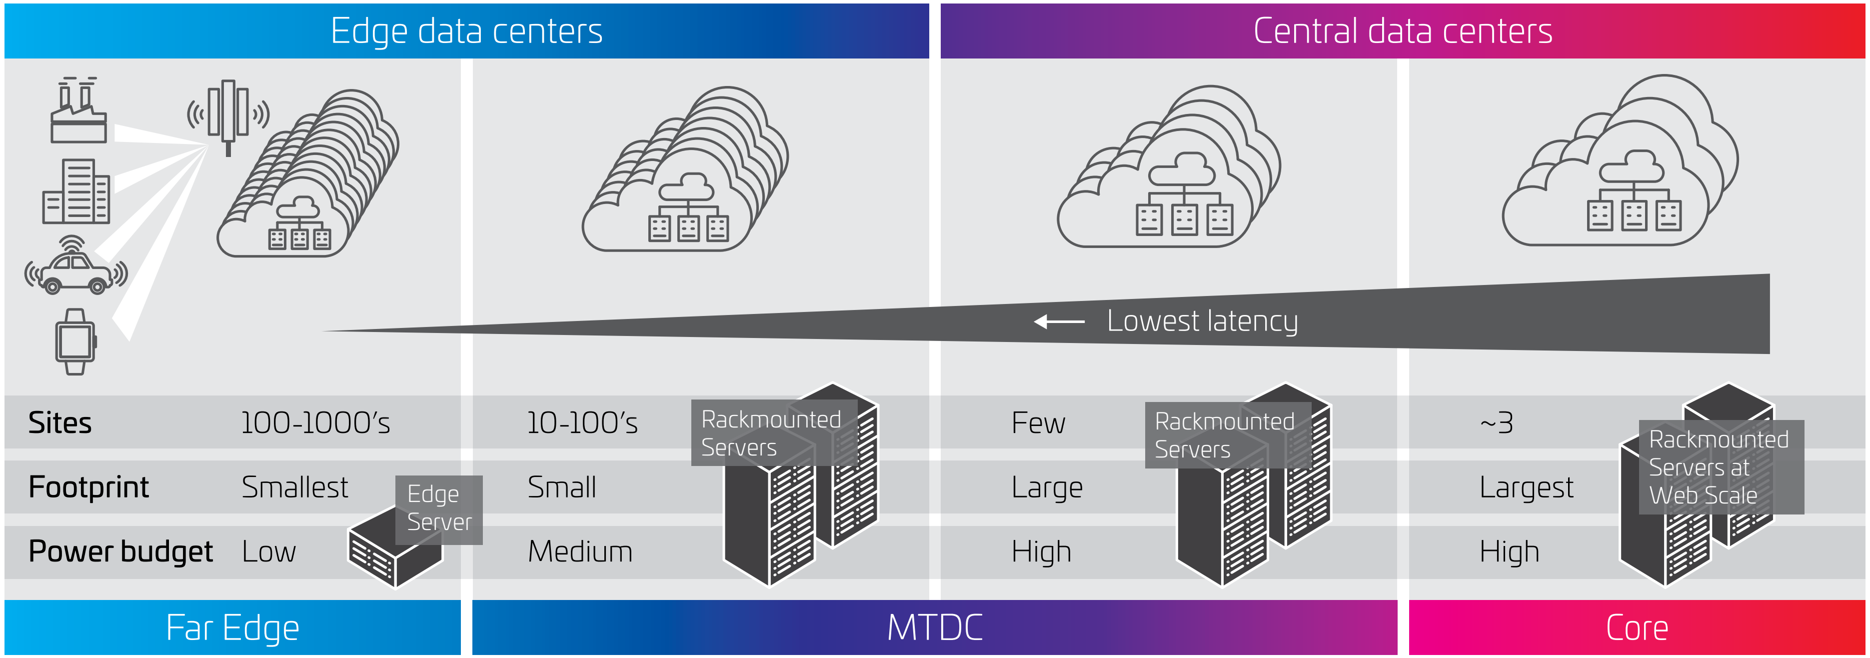

Figure 3 Source: Nokia, 2023

To ensure the operations at various MTDCs fit their intended purposes and can be interconnected as needed, a structured base build must be planned at the start of the project; this is especially true when external factors dictate siting the MTDCs across a campus, city, or region. A structured base build must be thought about from the start of the project. Standards provide valuable guidance when considering the base build process.

Standards such as TIA-942 and EN-50600 give solid guidance for data center infrastructure. They cover aspects of design, cooling systems, security and sustainability. Whilst not global standards, TIA-942 for North America, and EN50600 for Europe, they both provide a solid framework to ensure the MTDC meets the requirements for redundancy, availability and interconnectivity within the data center. Both standards specifically deal with the telecommunications cabling throughout the data center and the different key zones.

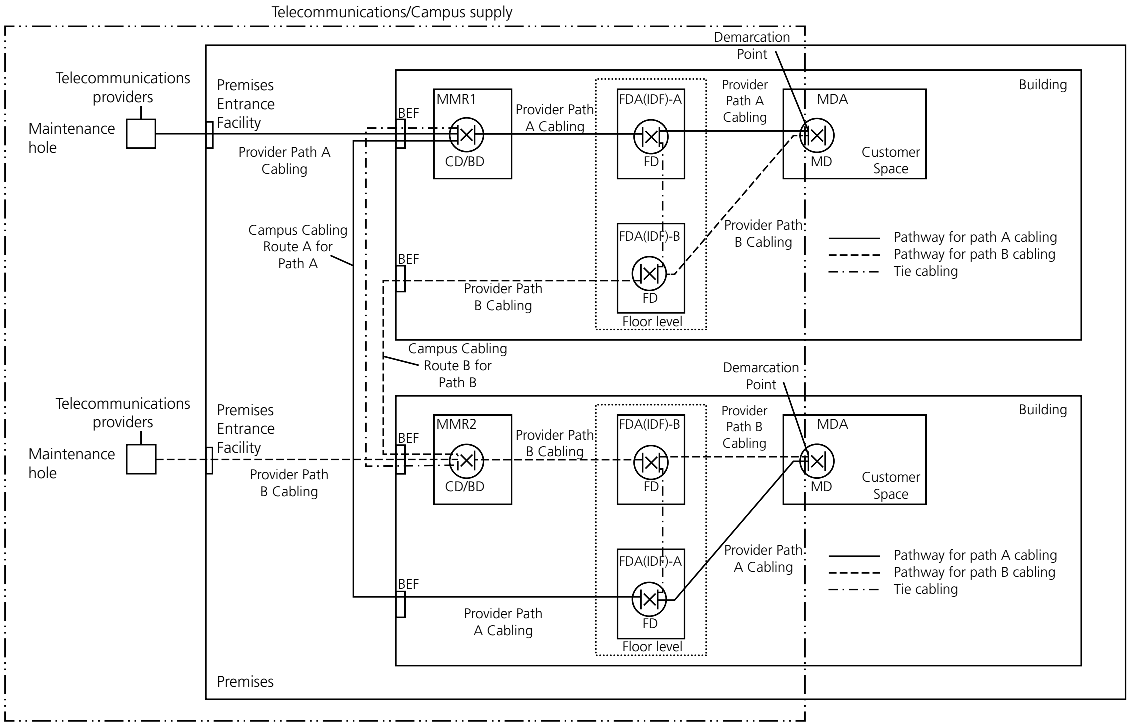

Figure 4: Key functional areas of the data center

As shown in Figure 4 (from left to right), the key functional areas of an MTDC are:

Building entrance facility (BEF) and premises entrance facility, which serve as transition points for external fiber and copper cables entering and exiting the MTDC. Transitions from the external cable to the internal cable are made with an approapriate optical fiber.

Meet-me rooms (MMRs) host the carrier equipment and cabling and distribute the cabling throughout the rest of the MTDC.

The floor distribution area / intermediate distribution frame (FDA/IDF) enables local floor-level flexibility between the MTDC structured cabling and the customer's cage.

The main distribution area (MDA), also known as a "demarc" or "demarcation point," is typically the last handover from the MTDC network to the customer's cage or hall. From here to the EDA (equipment distribution area), the customer will have control of the network and the operation will appear to be more like a typical data center operation. Note that all of the cabling paths are fully redundant, with A and B cabling paths.

Figure 5 shows some of the core infrastructure components and where they can be found across the various functional areas of the data center.

Fiber Entrance Cabinet (FEC)

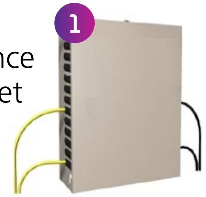

Rollable Ribbon Fiber Cable

High Density FO Patch Panels

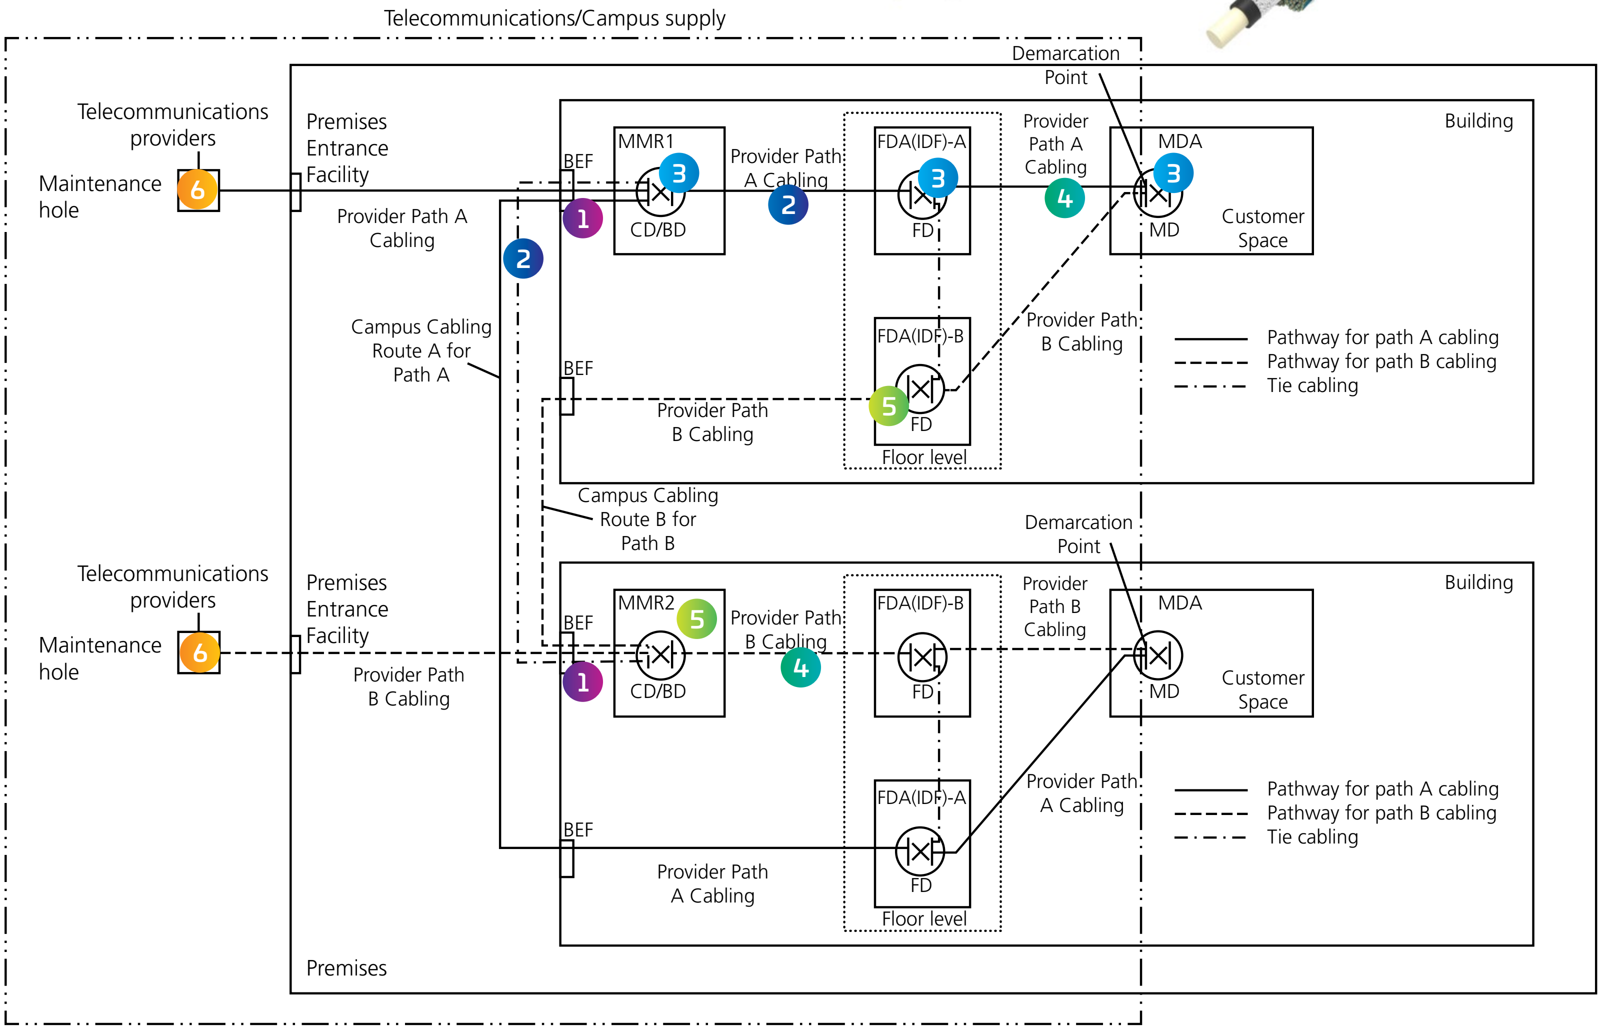

Figure 5

MPO Cable Assemblies

Optical Distribution Frame

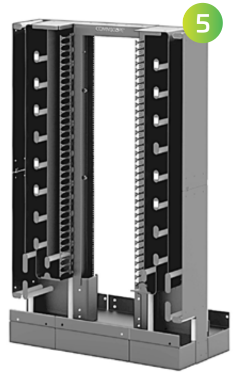

Fiber Optic Splice Closure

Conclusions

The three primary factors of power, government and population growth lead to a few conclusions that need to be considered when building an MTDC. The data center operator must design for expansion. Efficiency is the key.

Efficient power usage: As much grid energy as possible should be used to power the IT equipment instead of being wasted by inefficient hardware.

Efficient design: The base build should be flexible enough to be upgraded in support of new customers and data halls (wherever they are located). Look to the TIA942 and EN50600-2-4 standard for guidance.

Efficient supply chains: Explore how partners and supply chains can help build data centers in new locations. A holistic global approach to product selection can enable a partner to stage product sets—delivering them using a just-in-time approach to delivery. It can also simplify and speed installation in the field.

[i] Digital Realty Trust, Inc.; Q2 2022 Earnings Call, transcript; July 28, 2022

[ii] West London electricity capacity constraints; Mayor of London, London Assembly, briefing; July 2022

[iii] Cap on data centres ruled out despite surge in energy use; Irish Times, article; June 13, 2023

4

Data sovereignty

The tortoise and the hare

In the fable, “The tortoise and the hare,” a slow yet persistent tortoise challenges a hare to a race. The hare was far faster but also way overconfident. The slow but persistent tortoise kept moving along the path and, as we all know, won the race.

Currently, a similar race is playing out around the world as national governments slowly but surely progress toward their data sovereignty objectives. Until recently, these slow-moving bureaucratic institutions had a hard time writing legislation that would keep pace with more nimble private entities that were quickly expanding their digital services across international boundaries. That has now begun to change.

When the 27-member European Union (EU) adopted the General Data Protection Regulation (GDPR) in 2018, the new law initiated a wave of similar legislations that helped jump-start global efforts to ensure data sovereignty. The introduction of GDPR provided a legal framework that enabled federal governments to, for the first time, levy financial penalties on corporations that had previously operated outside national jurisdictions. The potential severity of the penalties was a wakeup call for CEOs and CIOs to take the issue of data sovereignty seriously. Simultaneously, legislators in other countries and local governments started to consider how they could use similar frameworks to protect their citizens and industries. However, national efforts such as these often come with unintended consequences.

National data sovereignty legislation

Figure 1

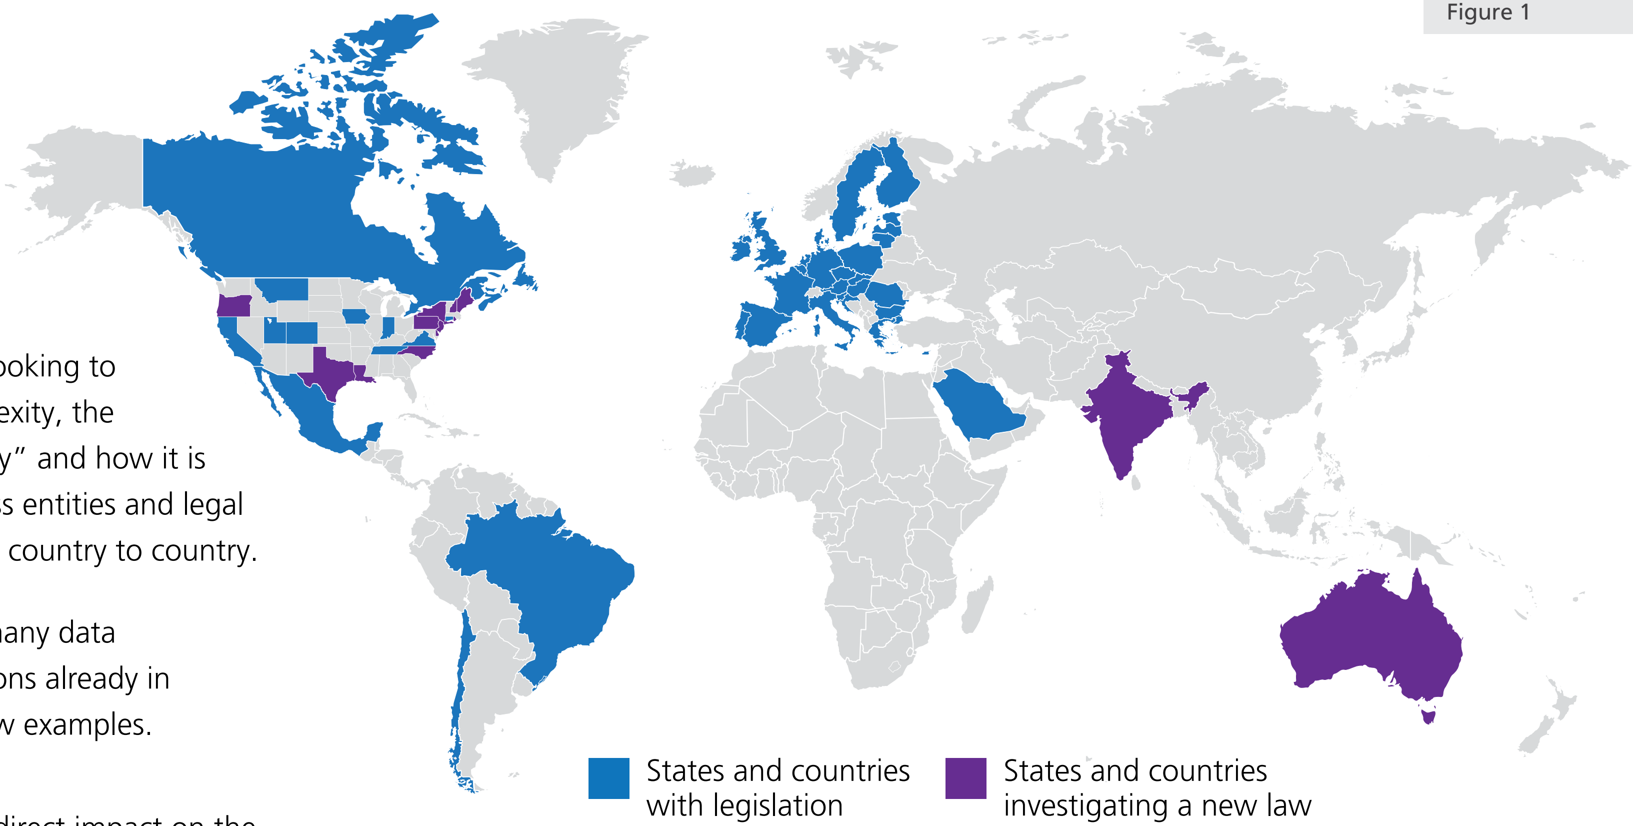

While large, multinational privacy regulations like GDPR make the headlines, there are countless smaller, regional laws that receive less attention but create key obstacles for multinational businesses looking to expand. Adding to the complexity, the definition of "data sovereignty" and how it is applied to individuals, business entities and legal transactions often varies from country to country.

Around the world there are many data sovereignty laws and regulations already in place. Figure 1 illustrates a few examples.

Each set of regulations has a direct impact on the data centers operating within the state, region or country. For example, the EU is currently considering a plan that will force providers to store all their data within the bloc and require a cloud cybersecurity certification. Further, ENISA, the EU's cybersecurity regulator, is drawing up new stricter requirements to ensure no foreign government can access EU data. Therefore, non-EU organizations may have to create "sovereign cloud" operations. Such cloud operations would need to be entirely located within the EU and compliant with EU rules that supersede all other national regulations. This could prove to be problematic to U.S. and Chinese tech companies. As an alternative, companies and EU-headquartered cloud providers are together considering partnerships that could provide a workaround—at least temporarily.

Data sovereignty vs data residency

At first glance, the terms "data sovereignty" and "data residency" may seem to be closely related or even interchangeable. They are not.

Data sovereignty refers to the laws and governmental policies that apply to data stored in the country where it originated and is currently located. In light of the increasing globalization of the world's data and the rapid adoption of cloud systems, it is easy to understand the difficulties of enforcing and operating within the various data sovereignty guidelines.

Data residency refers to the decision of businesses to store data outside of the jurisdiction where it was created. Once the data is moved (and made available for storage or processing) it is subject to the laws, customs and expectations of that specific region.

In summary, "data residency" refers to where the data is physically and geographically stored, while "data sovereignty" refers to the laws and governmental policies applicable to data stored in the country where it originated and is geographically located.

GDPR PERSONAL DATA

The EU's General Data Protection Regulation defines personal data as any information related to a person that can be used to directly or indirectly identify them, including:

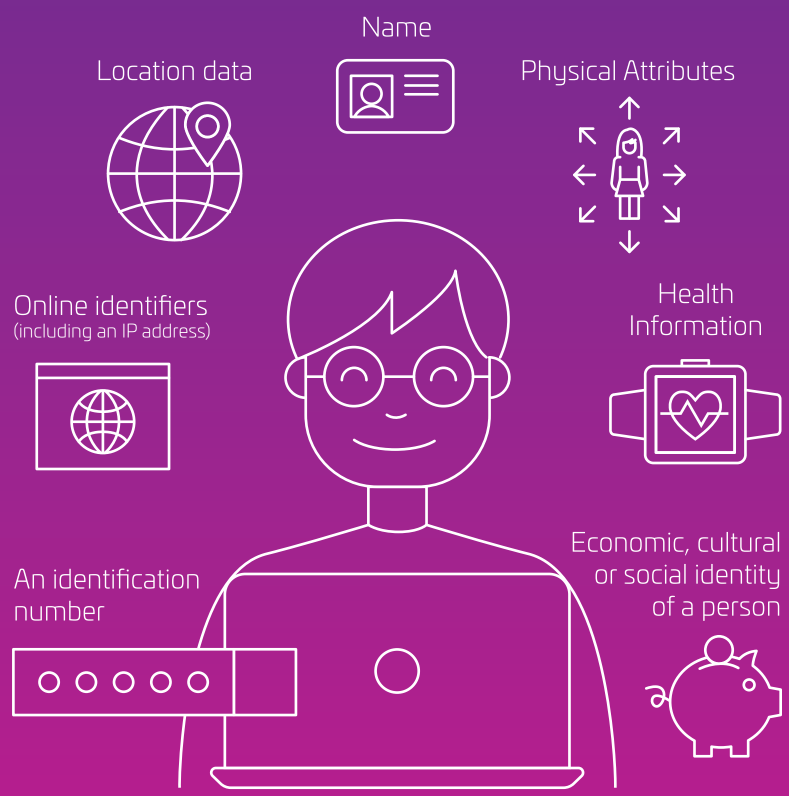

The challenges of a fractured system of data sovereignty

For large multinational businesses, navigating the ins and outs of each state's and country's data sovereignty regulations can be a significant impediment to global growth. Decisions as to where to locate a new facility can no longer be made purely on the strength of the business case; companies must also account for the local regulations and how they affect where data can be stored and processed. The following are just a few examples.

Data type and geography

Among the first considerations is the type of data to be stored and processed. For example, is it protected by personal privacy regulations (i.e., does it relate to a person's profile, employment, finance, health, and or payments)? Once the data type is characterized and understood, it must be evaluated within the context of the local or national data sovereignty laws. For example, some laws specify which type of data can, and cannot, leave the country of origin and cross national borders. Other laws allow moving some data types outside the country of origin, but only if the destination country has signed an equivalent privacy protection agreement (or law).

Finding the right data center fit

Depending on the data type and geography, there are four choices available for hosting the data storage and processing:

- **On-premises:** In this scenario, the data is processed and stored in an organization's own data center at a known location; this offers the best option for complying with the data sovereignty regulations. An on-premises data center can also be designed to match the agile cloud performance needed to support advanced applications like machine learning and AI. Today, more on-premises data centers are being built with a cloud-first approach.
- **Hybrid cloud:** The hybrid cloud blends cloud-based efficiencies from anything-as-a-service (XaaS) providers with localized on-premises resources. On one hand, it offers the flexibility, scalability and cost structure of a large cloud provider—perfect for handling non-regulated data. On the other hand, data that must comply with the local regulations can be stored on-premises, enabling businesses to better manage diverse data sovereignty requirements.
- **Private cloud:** A private cloud involves the use of a massive cloud-based infrastructure—none of which is owned by the end user. However, the cloud provider can dedicate portions of the underlying IT infrastructure to a single customer and ensure customer access is completely isolated. However, as with a hybrid cloud, the private cloud involves some tradeoffs. For example, having the IT infrastructure totally isolated provides the best opportunity to track and audit how the data is being stored and processed, yet there is no guarantee that the data in a private cloud will comply with national or regional data sovereignty laws.
- **Public cloud:** A public cloud consists of masses of common IT infrastructure—none of which is owned by the end user. Public cloud houses data in off-premises data centers anywhere in the world, so the location and ownership of the data become non-issues.

Power and location

Once the type of hosting has been decided, it's critical to understand if the power required is available to support a new installation or an extension to an existing installation. The base unit of any data center is the server and racks that house the servers. On-premises hosting typically involves four to 20 installed servers, each consuming approximately 1 kW, depending on the processing required. An installation of 100 racks in an on-premises data center with a power usage efficiency (PUE) of 1:2 could see a power draw of 1–5 MW.

For a cloud-scale deployment, the power draw becomes far greater as the maximum server density will be required to support all operational models being offered by the cloud service provider. In this instance it isn't unusual to have 25 servers per rack, with each rack using 20–80 kW of power and thousands of racks per location. Whether the data center is cloud scale or on-premises, the location of compute and storage resources are directly impacted by the availability of power as well as data latency performance.

Physical infrastructure

No matter where the data resides, the infrastructure must be built upon a strong yet agile passive cabling foundation. The physical layer infrastructure must be flexible enough to support the migration to higher data rates while satisfying the evolving requirements of the active equipment.

Fortunately, data center network topologies have evolved significantly—making it easier and more efficient to support more stable, flexible and future-ready deployments and applications. One of the major changes involves the migration from a three-tier approach (core, access and aggregation layers) to a Clos switching architecture, commonly termed "leaf-and-spine" (see Figure 2). This newer topology is based on an any-to-any connectivity approach that is ideally suited for today's high fiber density designs. Flatter with fewer "hops" between servers, the architecture can be easily expanded; the only real limitation to horizontal expansion is the number of ports on the spine switches. Since the network is flatter and faster, the physical layer in cabling should be ready to support day 1 transmission speeds and future data rates.

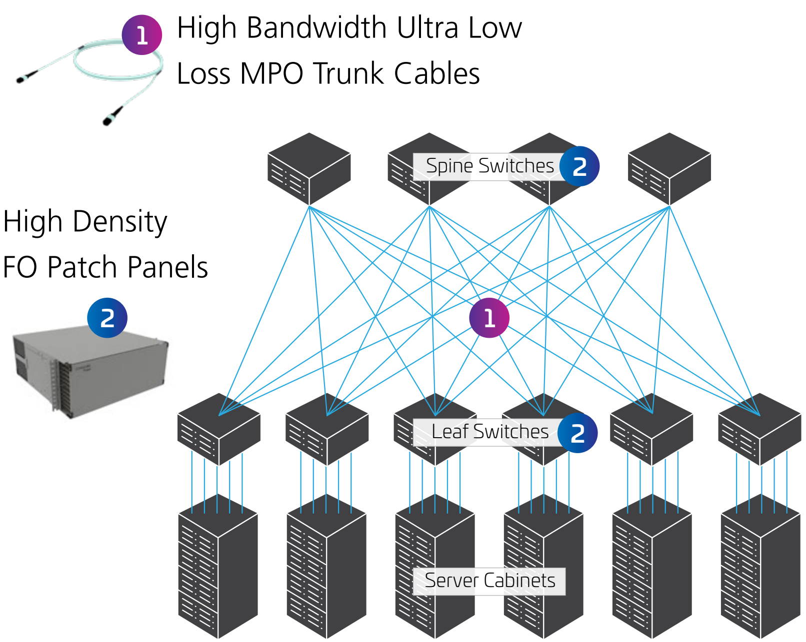

Figure 2: New flatter leaf-and-spine architecture using a three-tiered folded-Clos design for on-premises and hybrid data centers.

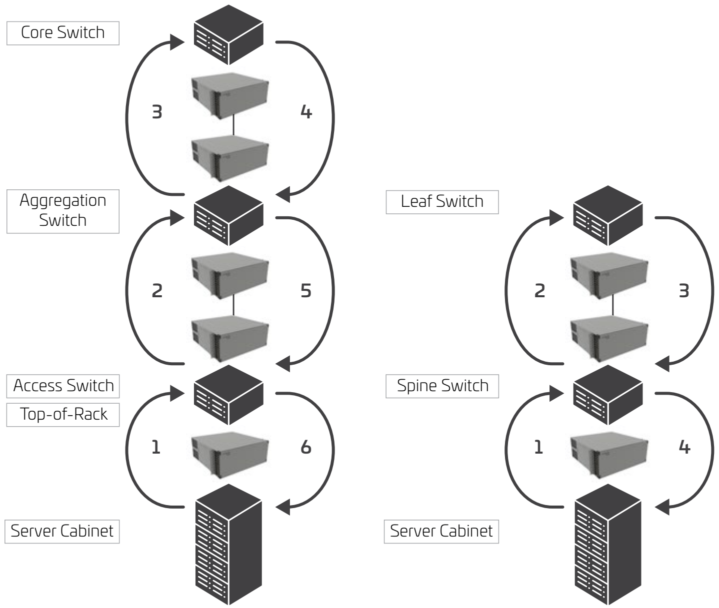

Figure 3: A traditional three-tier switching architecture versus a leaf-and-spine architecture

In Figure 3, the stack on the left uses a traditional three-tier switching architecture, requiring the data to make six separate hops to traverse the switching layers and reach the destination server. The right-hand stack illustrates the inherent benefits of the flatter leaf-and-spine architecture. Reducing the number of switching layers decreases the number of hops—and associated latency—up to 33 percent.

As adoption of cloud-scale architectures continues to ramp up, the industry is coalescing around the leaf-and-spine topology—with one small wrinkle: To satisfy the data handling requirements of much larger data centers, many network managers are adopting a three-tier leaf-and-spine topology, such as the one shown in Figure 4.

It is also important to note that, in the future, these architectures will increasingly be supported by 16-fiber MPO connectivity. As hyperscale and cloud-scale data centers migrate from 100G lane speeds to 400G, 800G, 1.6T and beyond, MPO-16 connectivity is the fundamental building block to higher speeds. Figure 4 shows how 16-fiber MPO connectivity supports the three-tier Clos network.

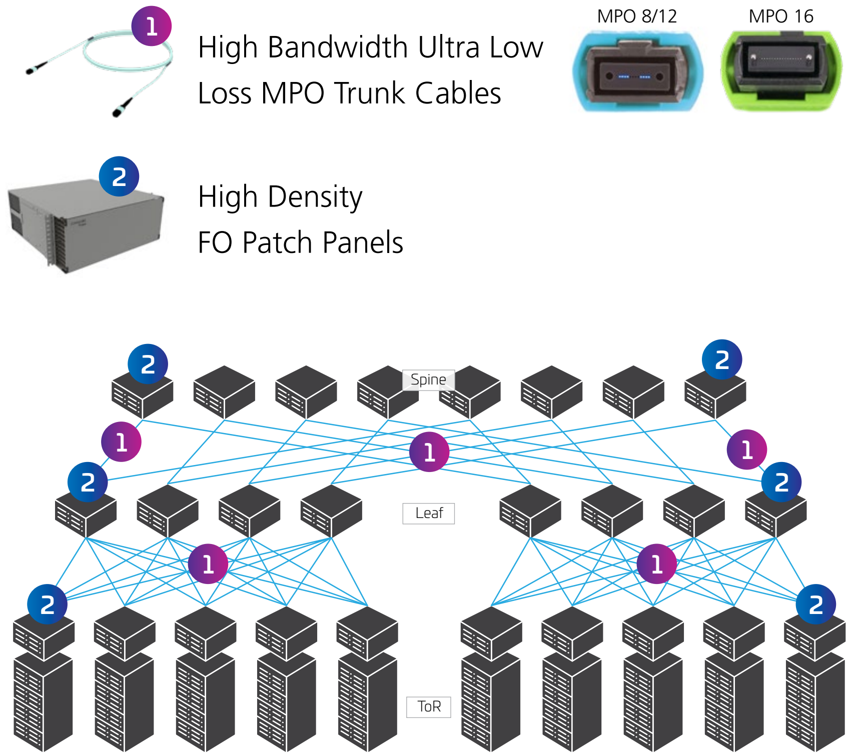

Figure 4: Three-tier Clos network with physical infrastructure components added, including MPO-16 connectors

Conclusion

While the definition of "data privacy" varies across regions, one thing that everyone agrees on is that it is critically important—particularly in an age of rapid network globalization. Looking into the future is always challenging, but if the past decade is any indicator, local variations in the generation and application of data sovereignty law are likely to increase.

Building a data center infrastructure able to support data sovereignty is essential but not impossible. Success hinges on being able to incorporate flexibility into the physical infrastructure to ensure it can support future topologies and data rates for the next generation of servers and switches.

With that in mind, the data center manager must be fluent in new connectivity technologies like MPO-16 and the fabric cabling needed to support growing bandwidth demands.

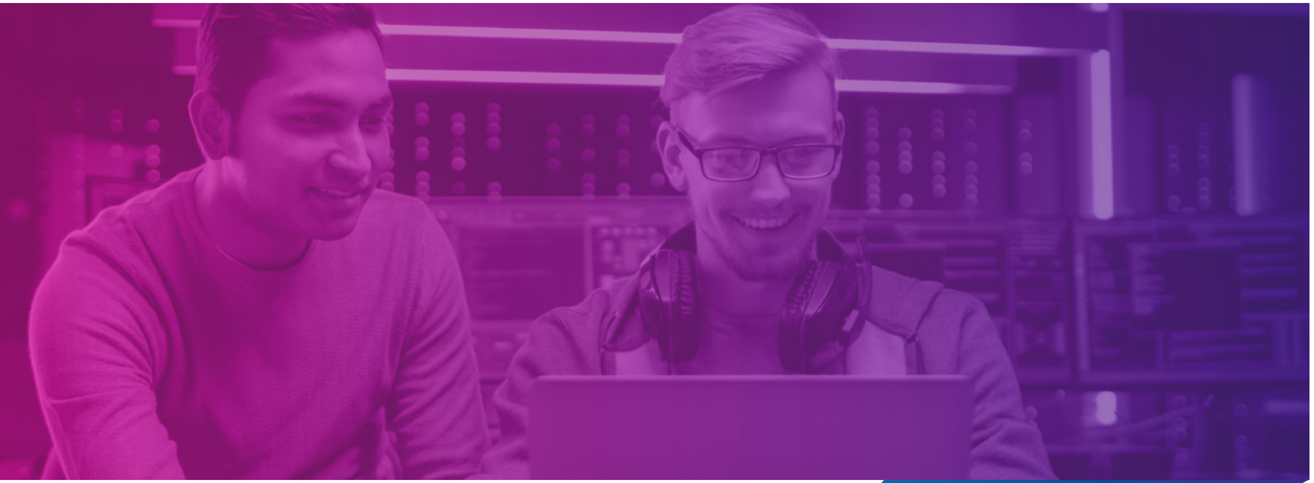

5

Sustainability in the data center

When we talk about sustainability, our discussions typically revolve around our carbon footprint and how the decisions we make today may impact our future. As an industry leader, CommScope is constantly asking ourselves hard questions:

Are we doing all we can to create a more sustainable industry and be a role model of corporate citizenship?

How can we further reduce or offset our environmental impact while continuing to improve how the world communicates and collaborates?

What new technologies and strategies will address the needs of our environment as well as our suppliers, partners and customers?

At the same time, the world is growing ever more data centric. Between 2010 and 2020, the volume of data generated, harvested, copied, and consumed worldwide grew by almost 5,000 percent—and data usage increased from 1.2 trillion gigabytes to 59 trillion gigabytes. [i] The exponential growth in data use shows no signs of slowing; the same is true for the new technologies and network infrastructure needed to support it.

Enabling the type of application connectivity, bandwidth and latency performance needed to keep pace with society's demands requires geographic diversity, capacity and accessibility on a massive scale. Service providers are responding by continuing to build out data centers of various types and sizes and investing in more interconnect networks to create more capacity and lower latency.

At first glance, the continued buildout of new ICT network infrastructure would appear to be in direct opposition to our goal of environmental sustainability. But this is not necessarily the case. In this chapter, we'll show how CommScope is successfully balancing the growth of data center capacity and the greening of our planet.

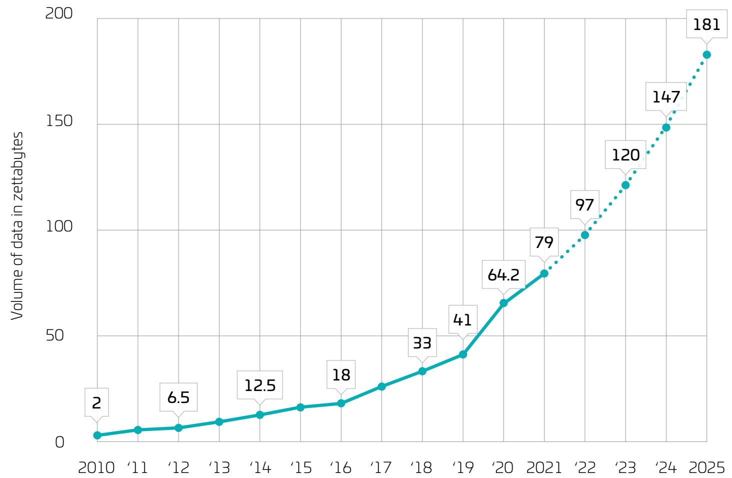

Figure 1: Global data generated, consumed, copied and stored
Source: Statista.com

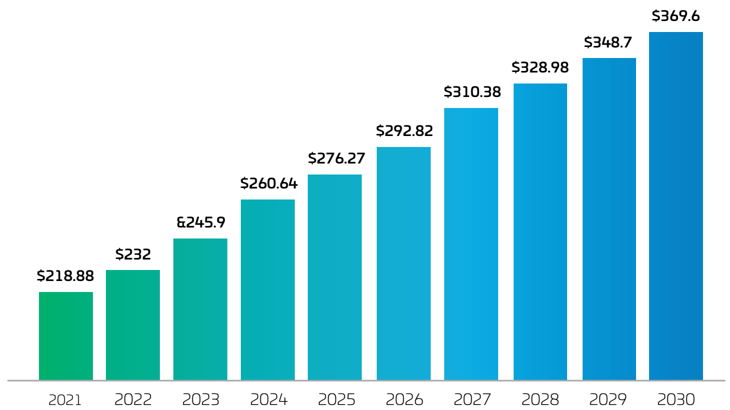

Figure 1: Global data generated, consumed, copied and stored
Source: Statista.com

Data centers rise to the environmental challenge

Despite its historically cautious mindset, the data center industry has taken decisive steps to address its environmental impact. 451 Research's report—Voice of the Enterprise: Datacenters, Sustainability 2023—shows the environment playing a very or somewhat important role in 76 percent of enterprise technology decisions, including data centers. [ii] More specifically, data center providers are focused on reducing the carbon intensity of the energy being consumed, with initiatives such as reducing water use and eliminating diesel generators as common topics of discussion. Other findings from the report show that:

89%

say data-center efficiency and sustainability are very or somewhat important.

80%

believe their data center's water consump-tion is a major or moderate concern.

36%

of companies that closed a data center did so to improve their carbon footprint.

In its executive summary of the report, S&P Global Market Intelligence, which owns 451 Research, concluded:

"The requirements for sustainability as well as changes in compute infrastructure (such as higher-density chips) will push the datacenter industry to test and adopt innovative technology. [These include] liquid cooling, micro-modular datacenters, alternatives to diesel generators and other technical changes in datacenter construction and operation. Customers will continue to push for new approaches, and this could see the industry become a leader in sustainability and an example for enterprises more broadly."

Figure 3: Role of the environment on enterprise tech decisions
Source: 451 Research

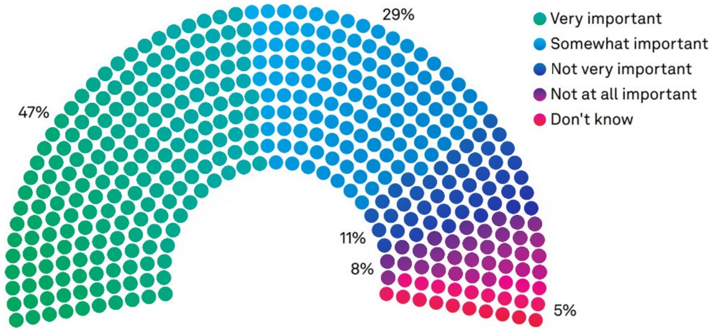

If data center operators are to successfully thread the needle between improved sustainability and increased capacity and performance, they'll need help from their infrastructure technology partners—and they're getting it in a variety of forms.

New fiber-based products and portfolios can now support multiple generations of application and equipment upgrades while reducing the amount of non-recyclable packaging and installation labor. Advanced cabling configurations are also helping simplify network architectures, creating a positive impact on space and energy consumption. New combinations of passive and active components are now demonstrating the potential to maximize application performance and access, benefiting customers and the environment.

Going faster while reducing resource consumption

Over the past few years, data rates have increased rapidly, progressing beyond legacy 10G. Transceiver speeds are rapidly moving from 100G/200G to 400G, 800G, 1.6T and beyond.[iii]

By 2027, shipments of 100G servers are expected to dominate the market. To support a typical 1:1 subscription ratio, the leaf-and-spine fabric needs to run at 400G.

In the past, faster chipsets meant an increase in power consumption and cost. However, thanks to recent advancements in parallel optics and breakout options, a single high-speed switch port can now support one, four or eight different devices. This means fewer switches are required to provide the same or more capacity—with one switch now doing the work of six. Moreover, these new capabilities reduce the number of switch layers, lower the cost and power per gigabit, and reduce network complexity and the amount of mined and processed minerals and manufacturing. Those efficiencies are driving data center operators to upgrade and migrate in advance of the traditional three-year refresh cycle.

Propel™—CommScope's end-to-end, ultra-low loss, modular structured cabling fiber platform—enables 1:1 matching of module and adapter options for application-based scaling, making it the most efficient fiber solution available.

The portfolio is also uniquely sustainable, designed to reduce waste, conserve space, reduce fuel use and extend the product lifecycle:

- 4U panel packaging is 20 percent smaller and 16 percent lighter than typical panels
- Installs with one technician, reducing CO2e.
- Supports multiple upgrades, reducing mined materials and manufacturing impact
- Components use conflict-free minerals

MPO16 connectivity takes sustainability further:

- 62 percent less plastic (fewer trunks, more links)
- 61 percent less steel (ultra-high-density panel)
- 57 percent less cardboard, 33 percent less packaging

The trend to 8- and 16-fiber connectivity

Historically, 12- and 24-fiber-based units were the norm for fiber cable construction. However, the move to faster lane speeds has fueled a migration to 16-fiber MPO connectivity, the basic building block for speeds of 400G and higher. 16-fiber as well as 8-fiber connectivity have led to more application-specific fiber configurations and modular connectivity, which have simplified everything from design and installation to Day 2 operations.

Going forward, it is expected that data center applications will be delivered over 2, 8 or 16 fibers. This means that, to maximize switch port utilization (and minimize power consumption), the trunk cabling must align with the applications being supported. Duplex applications can be supported by all trunk options, while 8- or 16-fiber applications are best supported by 8- or 16-fiber trunks, with the 16-fiber option providing the most flexibility. In lieu of 16-fiber trunk cables, combining equal lengths of 8-fiber trunks as needed can also be done to support 16-fiber applications.

Choosing the right building block(s) for data center connectivity can go a long way toward helping data centers meet their sustainability objectives. As mentioned, the 16-fiber trunk is the most flexible and efficient way to support speeds of 400G and above. Therefore, network managers can reduce their power draw and get the biggest bang for their cabling buck. The 16-fiber option also enables greater efficiency in onsite labor and space requirements. Thus, it can improve Day 2 operation efficiency, reducing the need for truck rolls and the associated generation of CO2 emissions.

Rethinking fiber cabling construction

Yet another way in which data center operators are improving their environmental outcomes is with new and innovative fiber cabling designs. One of the best examples is rollable ribbon fiber cable. While smaller diameter cables typically weigh less and provide pathway weight and space benefits, the development of rollable ribbon fiber cable has enabled much higher fiber counts in much smaller jacket diameters.

Rollable ribbon fiber cable is based, in part, on the earlier development of the central tube matrix ribbon cable, which featured ribbon stacks of up to 864 fibers within a single, central buffer tube. The fibers are grouped and continuously bonded down the length of the cable, which increases its rigidity. While this affects little when deploying the cable in an OSP application in a data center, a rigid cable is undesirable because of the limited routing restrictions these cables require.

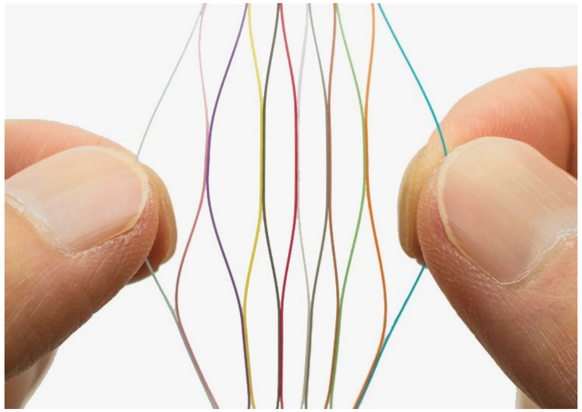

Figure 4: Intermittent bonding of rollable ribbon fibers

In the rollable ribbon fiber cable, the fibers are attached intermittently to form a loose web. This configuration makes the ribbon more flexible, allowing the fibers to flex with a degree of independence to one another. The fibers can now be “rolled” into a cylinder, making much better use of space when compared with flat ribbons.

As many as six 3,456 rollable ribbon fiber cables can be loaded into one four-inch duct, more than twice the density of conventionally packed fibers. Compared to traditional matrix ribbon cables, rollable ribbon cables use significantly less plastics and related materials while delivering more fiber protection and routing flexibility. The design also leads to improved on-site efficiencies during installations.

Low-impact packaging

Needless to say, the type and design of the packaging of fiber cabling—or any infrastructure component—will affect the data center's carbon footprint as well. The challenge is how to best protect the product while minimizing the environmental impact of the packaging after installation. One way to do this is to use recycled and/or recyclable packaging materials versus single-use plastics, wherever possible. Developing green-focused materials along with advances in recycling have enabled much more efficient and environmentally beneficial options.

Another consideration with regard to product packaging is its overall weight and size, as these factors figure prominently in the amount of fueled consumed and CO2 emitted during transportation. Careful consideration of those materials can impact usability and efficiency on site for the installers and data center operators.

Geographic production and supplier diversity also play a role. Supply chain constraints recently highlighted that as well. Sustainability planning considers availability of the resource today and in the future. It's a small world, after all. Diversity of pre-qualified component supply options is important in ensuring resources are not depleted environmentally and also, as we are witnessing in the capacity build, that they are available when and where needed.

Case in point

Recently, a large social media provider realized significant improvements in their carbon footprint by upgrading to a next-generation infrastructure architecture. They were able to reduce plastics, steel, cardboard and packaging materials, while simplifying their architecture and enabling at least three generations of connectivity.

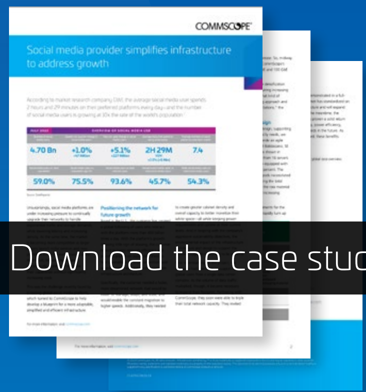

Download the case study

[i] 54 Predictions About The State Of Data In 2021; Forbes, article; December 30, 2020

[ii] 2023 Trends in Datacenter Services & Infrastructure; 451 Research, report; December 2022

[iii] New transceivers that will use MPO16: 800G-DR8, 800G-DR8-2, 800G-SR8, 800G-VR8, 1.6T-SR8.2, 1.6T-VR8.2, 1.6T-DR8, 1.6T-DR8-2

6

The evolving role of the Data Center in a 5G-enabled world

For decades, the data center has stood at or near the center of the network. For enterprises, telco carriers, and cable operators—and, more recently, service providers like Google and Facebook—the data center was the heart and muscle of IT.

The emergence of the cloud has emphasized the central importance of the modern data center. But listen closely and you'll hear the rumblings of change.

As networks plan for migration to 5G and IoT, IT managers are focusing on the edge and the increasing need to locate more capacity and processing power closer to the end users. As they do, they are re-evaluating the role of their data centers.

According to Gartner[1], by 2025, 75 percent of enterprise-generated data will be created and processed at the edge—up from just 10 percent in 2018.

At the same time, the volume of data is getting ready to hit another gear. A single autonomous car will churn out an average of 4T of data per hour of driving.

Networks are now scrambling to figure out how best to support huge increases in edge-based traffic volume as well as the demand for single-digital latency performance, without torpedoing the investment in their existing data centers.

A heavy investment in east-west network links and peer-to-peer redundant nodes is part of the answer, as is building more processing power where the data is created. But what about the data centers? What role will they play?

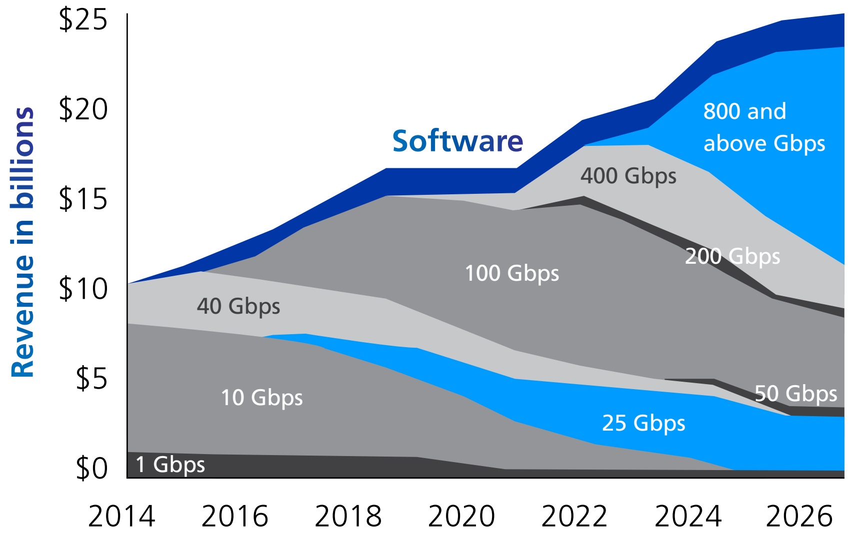

Source: 650 Group, Market Intelligence Report December 2020

1 What Edge Computing Means for Infrastructure and Operations Leaders; Smarter with Gartner; October 3, 2018

The AI/ML feedback loop

The future business case for hyperscale and cloud-scale data centers lies in their massive processing and storage capacity. As activity heats up on the edge, the data center's power will be needed to create the algorithms that enable the data to be processed. In an IoT-empowered world, the importance of AI and ML cannot be understated. Neither can the role of the data center in making it happen.

Producing the algorithms needed to drive AI and ML requires massive amounts of data processing. Core data centers have begun deploying larger CPUs teamed with tensor processing units (TPUs) or other specialty hardware. In addition, the effort requires very high-speed, high-capacity networks featuring an advanced switch layer feeding banks of servers—all working on the same problem. AI and ML models are the product of this intensive effort.

On the other end of the process, the AI and ML models need to be located where they can have the greatest business impact. For enterprise AI applications like facial recognition, for example, the ultra-low latency requirements dictate they be deployed locally, not at the core. But the models must also be adjusted periodically, so the data collected at the edge is then fed back to the data center in order to update and refine the algorithms.

Playing in the sandbox or owning it?

The AI/ML feedback loop is one example of how data centers will need to work to support a more expansive and diverse network ecosystem—not dominate it. For the largest players in the hyperscale data center space, adapting to a more distributed, collaborative environment will not come easily. They want to make sure that, if you're doing AI or ML or accessing the edge, you're going to do it on their platform, but not necessarily in their facilities.

Providers like AWS, Microsoft and Google are now pushing racks of capacity into customer locations—including private data centers, central offices and on-premises within the enterprise. This enables customers to build and run cloud-based applications from their facilities, using the provider's platform. Because these platforms are also imbedded in many of the carriers' systems, the customer can also run their applications anywhere the carrier has a presence. This model, still in its infancy, provides more flexibility for the customer while enabling the providers to control and stake a claim at the edge.

Meanwhile, other models hint at a more open and inclusive approach. Edge data center manufacturers are designing hosted data centers with standardized compute, storage and networking resources. Smaller customers—a gaming company, for example—can rent a virtual machine to host their customers and the data center operator will charge you on a revenue sharing model. For a small business competing for access to the edge, this is an attractive model (maybe the only way for them to compete).

Foundational challenges

As the vision for next-generation networks comes into focus, the industry must confront the challenges of implementation. Within the data center, we know what that looks like: Server connections will go from 50G per lane to 100G; switching bandwidth will increase to 25.6T; and migration to 100G technology will take us to 800G pluggable modules.

Less clear is how we design the infrastructure from the core to the edge—specifically, how we execute the DCI architectures and metro and long-haul links, and support the high-redundancy peer-to-peer edge nodes. The other challenge is developing the orchestration and automation capabilities needed to manage and route the massive amounts of traffic. These issues are front and center as the industry moves toward a 5G/ IoT-enabled network.

Getting there together

What we do know for sure is that the job of building and implementing next-generation networks will involve a coordinated effort.

The data center—whose ability to deliver low- cost, high-volume compute and storage cannot be duplicated at the edge—will certainly have a role to play. But, as responsibilities within the network become more distributed, the data center's job will be subordinate to that of the larger ecosystem.

Tying it all together will be a faster, more reliable physical layer, beginning at the core and extending to the furthest edges of the network. It will be this cabling and connectivity platform—powered by traditional Ethernet optics and coherent processing technologies—that will fuel capacity. New switches featuring co-packaged optics and silicon photonics will drive more network efficiencies. And, of course more fiber everywhere—packaged in ultra-high-count, compact cabling—that will underpin the network performance evolution.

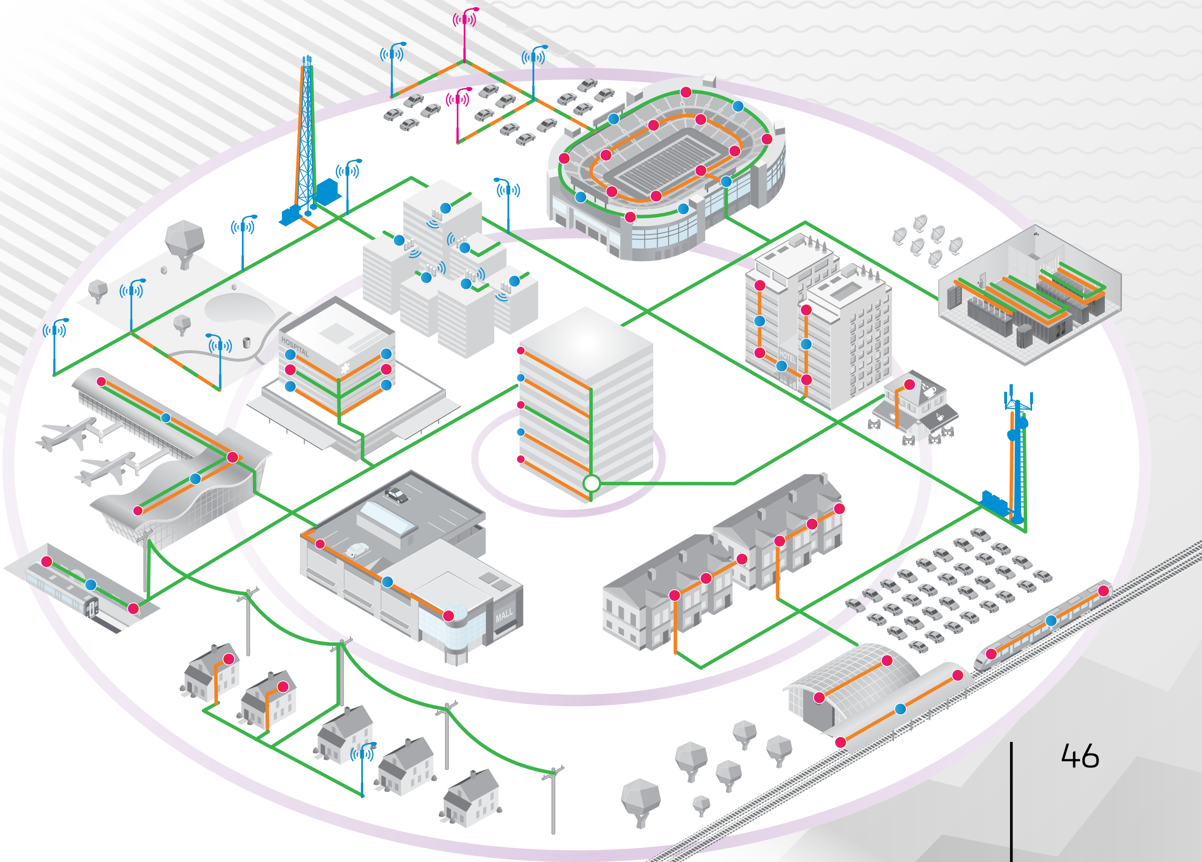

7

Data center interconnect (DCI)

Across the campus and into the cloud

Digitization is driving more traffic to more data centers across the network and between enterprise data centers, disaster recovery locations, and eventually to multiple cloud peering points. To maintain good quality of service, traffic needs to be sent across the campus and out into the cloud rapidly and securely. Therefore, Data Center Interconnect (DCI) has become essential, as data centers no longer exist as individual and isolated islands but are now part of a highly interconnected ecosystem.

Driven by these and other applications, the DCI market, between 2023 and 2028, is predicted to grow by over 16 percent CAGR [i].

What is a data center interconnect?

A data center interconnect refers to the infrastructure that connects two or more individual data centers involved in a common task. The geographic scope of a DCI varies greatly. Multiple data centers within a campus may define a campus-scale DCI while clusters of data centers across several cities may constitute a regional DCI (also known as an "availability zone"). At the large level, data center networks spread across the world are connected to create a global DCI.

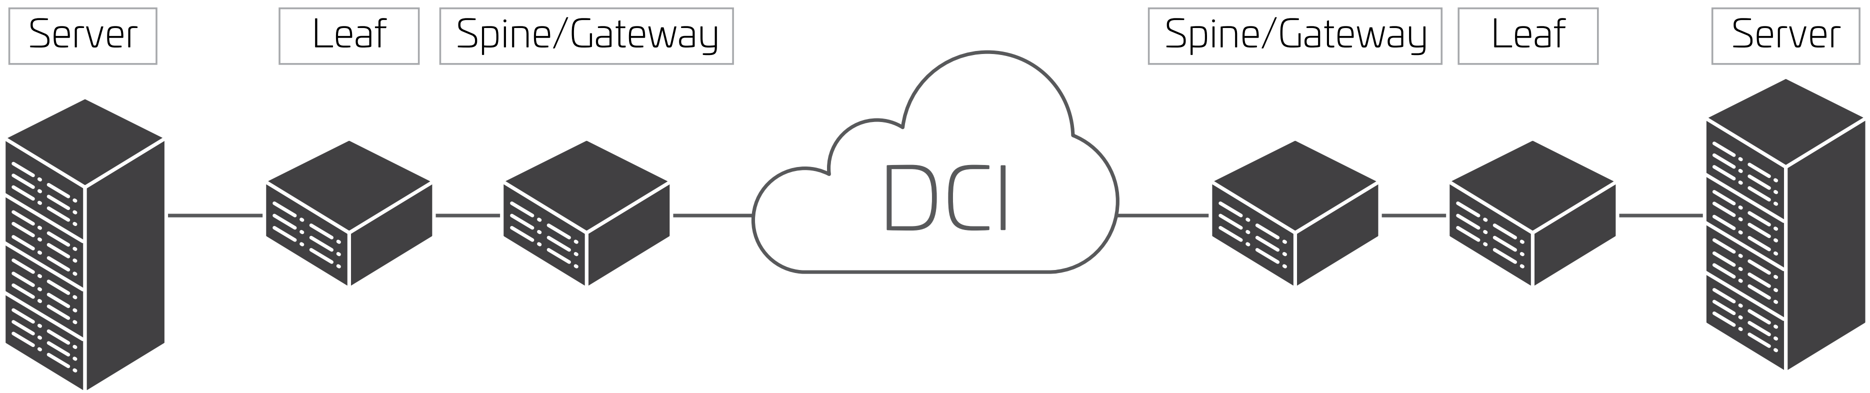

Figure 1: Basic architecture of a DCI

DCI transmission

There are any number of ways to transmit data between data centers in a DCI. The technology can involve sending high-speed Ethernet or optical signals over dedicated fiber or wavelength services. In most cases, a DCI requires a high-speed WAN link, which could mean using MPLS, Ethernet, VPLS, metro Ethernet, etc.

Selecting the right DCI infrastructure, architecture and topology depends on a wide range of variables, including the location of the data centers, the distance between data centers, bandwidth and availability requirements, the capabilities of local service providers, and security concerns.

There are just as many options and variables when you get down to the component level. Here are just a few of the developments that may figure into your decision.

Border edge (gateway) devices

Leaf-and-spine switches are integral to creating a data center fabric and are used to ensure efficient east-west transmission of data. As a data center approaches cloud scale, the spine-layer switches are typically dedicated to facilitating the east-west data flow. A new group of switches, known as "border edge switches," has been added to the architecture to handle incoming and outgoing traffic at the edge of the data center fabric. In some cases, the job of handing off data to the DCI network can be performed at the leaf layer, using a border leaf switch.

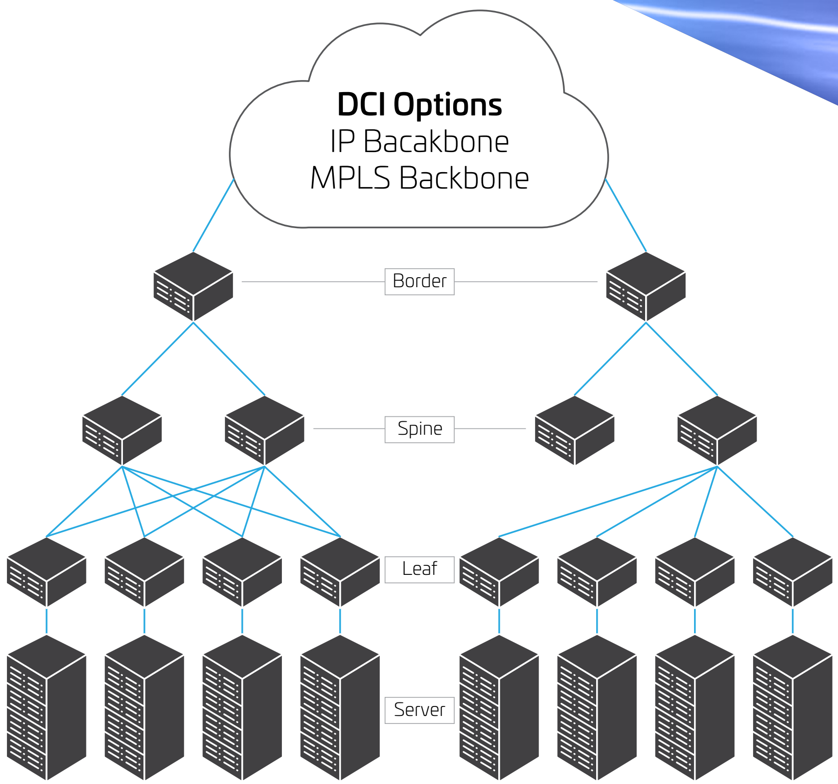

Figure 2: Placement of border-edge switches

Coherent optics

Coherent transmission has been used in long-haul and undersea links for about a decade, but we now see this technology migrating to the data center. Over time, the makers of coherent transceivers have reduced the size, power and cost of their optics to be more attractive for shorter and shorter links. IEEE Ethernet standardized the use of 100G and 400G coherent signaling on fiber links up to 80 km. A new project (802.3dj) will write standards for 800G coherent over just 10 km.

Figure 3: QSFP-DD coherent optic module

Figure 4: CFP2 coherent optic module

As data centers continue to grow and become more interconnected, the bandwidth needs of the DCI network are also growing, reaching 100 Tbps over multiple wavelengths. This bandwidth demand is supported by faster data rates—400 Gbps and 800 Gbps per wavelength—and will only continue to escalate.

Enabled by the IEEE802.3ct standard, coherent optics are typically used for ultra-high bandwidth applications ranging anywhere from 100G to 1T over very long distances. Powerful digital signal processing chips (DSPs) are embedded within these systems to mitigate linear effects caused by fiber impairments, including chromatic dispersion and polarization mode dispersion.

Coherent fiber optics utilize the natural properties of light to optimize digital modulation practices and fiber-optic carrying capacity in long-range applications. However, coherent transmission will change to suit DCI applications. Supported by a proposed new standard (IEEE802.3dj), this will attempt to enable 800G links over just 10 km.

Dense wave division multiplexing (DWDM)

DWDM enables multiple wavelengths of light to travel over the same fiber simultaneously, with each wavelength carrying a discrete signal. Tight wavelength spacing can enable up to 96 channels on a single fiber. DWDM is a versatile transmission technology able to support coherent optics and on-off key (ook) amplitude signaling. When combined with coherent modulation, individual channel bandwidth can expand to 400 or 800 gigabits.

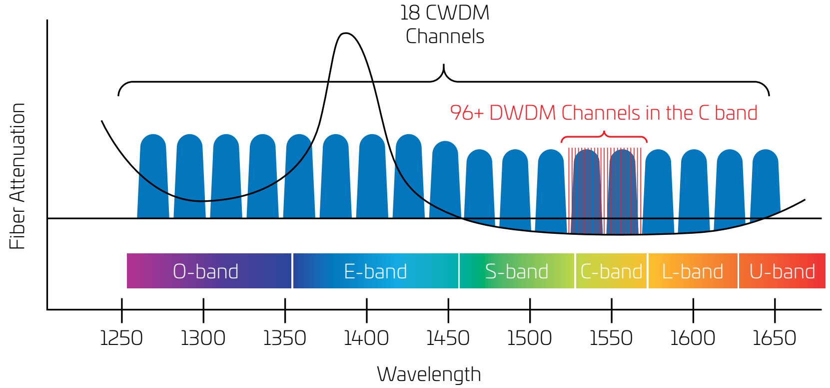

Figure 5: CWDM/DWDM Band and Channel Configurations

QPSK coherent coding vs PAM4

Coherent systems are based around phase shift keying (PSK)—phase modulation techniques that allow multiple symbols per bit to be encoded based on four phase shift orientations (e.g., 0°, 90°, 180°, and 270°). Many systems use quadrature phase shift keying (QPSK) to encode two bits per symbol. Dual polarization QPSK (DP QPSK) uses horizontal and vertical polarization along with QPSK to represent twice as many bits.

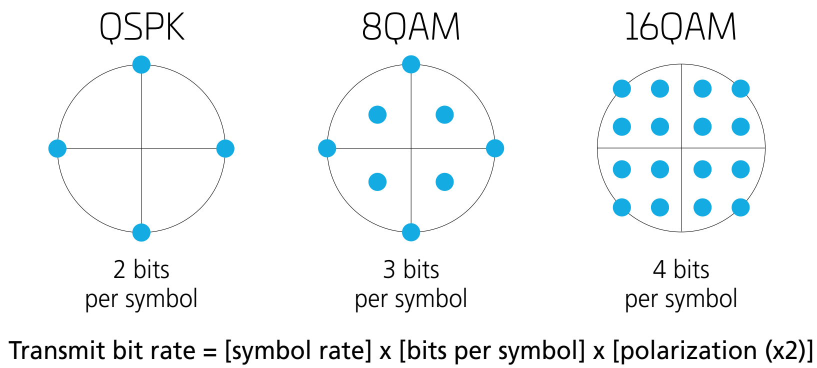

Figure 6: High-order modulation—Constellation™ diagrams

Pulse amplitude modulation (PAM4) is a four-level modulation scheme designed for short-haul fiber links. PAM4 (see figure 7) uses four amplitude pulses, each containing two bits, to double the bandwidth of conventional binary signaling. The simplicity and low power requirements of PAM4 make it a popular option for 100G and 400G Ethernet applications.

Unlike coherent optics, PAM4 is highly susceptible to fiber impairments. This limits range to ≤30 kilometers whereas coherent optical fiber communication systems can potentially span thousands of kilometers on amplified links.

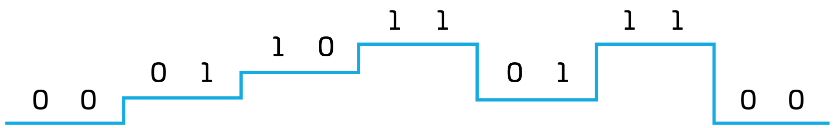

Figure 7: PAM4 signaling technology

New designs driving higher fiber-count DCI networks

The expansion of east-west data flows and the move to a low-latency leaf-and-spine switching architecture has created a tsunami of data inside and among data centers. This has pushed the development of data center campuses with multiple buildings into overdrive. This trend isn't isolated to a few hyperscalers in the U.S. It also impacts cloud and MTDC providers globally.

Once the cable has been terminated inside the data center, it needs to be presented for splicing to an onward destination or connected to patching and cross-connection equipment. This is done using optical distribution frames (ODFs) located in the meet-me room (MMR) or main distribution area (MDA), where all network cabling comes together and is prepared for distribution.

To ensure this zone can support future growth and is manageable, patch cord hygiene should be mandatory. As noted earlier, the role of the DCI network across the campus is increasing; in the near future, they will have to support throughput of 100T or more. This will require thousands of fibers—all converging at the MMR or MDA, thus the importance of patch cord management. So, too, ensuring the ODF has both fiber patch cord routing and slack management will be key to ensuring the MDA and MMR can support all future growth needs.

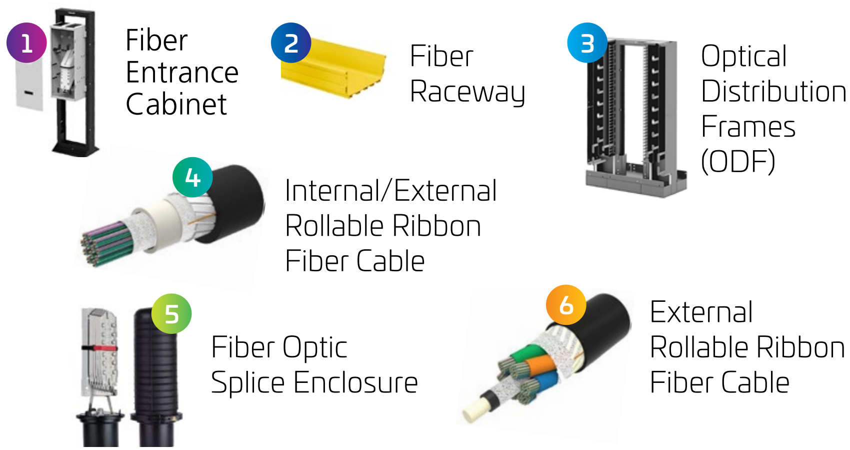

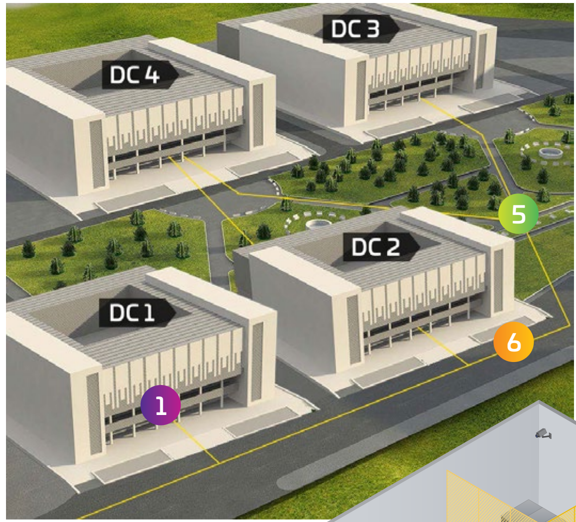

Figure 9: Infrastructure components for a typical campus-scale DCI

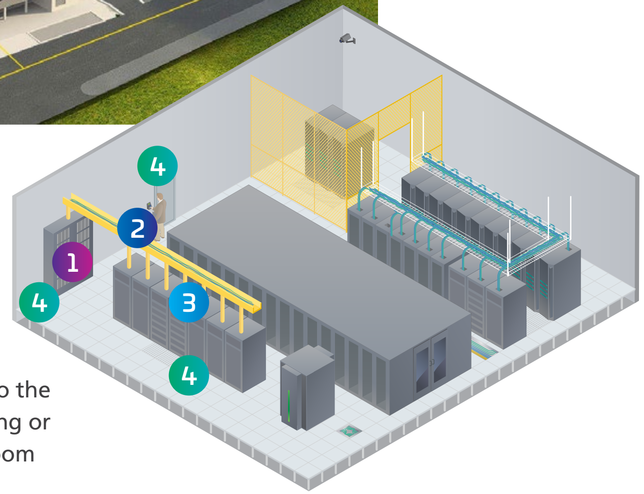

Figure 8: Entry into the data center building or MTDC meet-me room

To the future

Increasing deployment and expansion of DCI networks across the campus and into the cloud will continue unabated. As switch architectures flatten in support of machine learning and artificial intelligence, more and more data is being created and shared, mostly by machines that—unlike humans—never take a break or grow tired. As a result, data loads and DCI network requirements keep increasing. This presents physical challenges for the campus network designer, especially if the network has to interconnect data centers across a metro area or land owned by others.

If the data is to traverse these geographies seamlessly, the data-carrying capacity of the fiber cables must increase. This will be done either by adding more fibers per cable—via size-reduced cladding—or introducing a commercially viable alternative glass technology.

[i] Data Center Interconnect Market to Grow at CAGR of 16.03% through 2028; The Brainy Insights, news release, April 5, 2023

8

Cabling considerations for AI data centers

Introduction

For decades the danger of malicious artificial intelligence (AI) has been a trope in science fiction. Film antagonists like HAL 9000, the Terminator, the Replicants and the robots from The Matrix are opposing forces to the plucky humans who must overcome the dangers of technology. Recently, the release of DALLE-2 and ChatGPT has captured the imagination of the wider public as to what AI can do. This has led to discussions on how AI will change the nature of education and work. AI is also the main driver for current and future data center growth.

There are three aspects to AI:

- During training, a large set of data is fed into the algorithm that consumes the data and "learns" from it.
- The algorithm is then exposed to a new data set and tasked with deriving new knowledge or conclusions based on what it learned during training. For example, is this a picture of a cat? This process is known as "inference AI."
- The third (and perhaps most exciting) aspect is what's known as "generative AI." Generative AI is when the algorithm "creates" original output—text, images, videos, code, etc.—from simple prompts.

AI computation is handled by graphical processing units (GPUs): specialized chips designed for parallel processing and well suited to AI. The models used to train and run AI consume a significant amount of processing capacity—typically too much for a single machine.

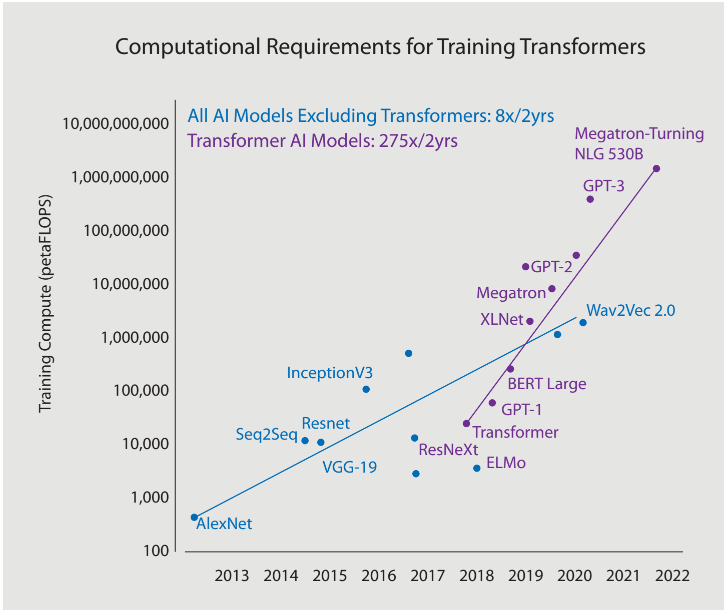

Fig. 1: AI model size in petaFLOPS
(source: https://blogs.nvidia.com/blog/2022/03/25/what-is-a-transformer-model/)

Figure 1 shows the historical growth of AI models in petaFLOPS (quadrillions of floating-point operations per second). Processing these large models requires multiple interconnected GPUs spread over many servers and racks. An AI data center deploys dozens of these AI clusters, and the cabling infrastructure that ties everything together to keep the data flowing presents a unique set of challenges.

The following outlines some of the key challenges and opportunities of cabling AI data centers, along with a few best practices and tips for success.

Typical data hall architecture

Nearly all modern data centers, especially hyperscale, use a folded Clos architecture, also called "leaf-and-spine." All the leaf switches in a data center connect to all the spine switches. In the data hall, server racks connect to a top-of-rack (ToR) switch. The ToR is then connected to a leaf switch at the end of the row or in another room via fiber cable. The servers in the rack are connected to the ToR with short copper cables—one to two meters long—carrying 25G or 50G signaling.

This configuration uses few fiber cables in the data hall. For example, Meta data centers that use the F16 architecture (see Figure 2) will have 16 duplex fiber cables from each of the server racks in a row. These cables run from the ToR to the end of the row, where they connect with modules that combine duplex fibers to 24-fiber cables. The 24-fiber cables then run to another room to connect to leaf switches.

Data centers that implement AI will house AI clusters next to compute clusters with traditional architecture. Traditional compute is sometimes called the "front-end network," and the AI clusters are sometimes called the "back-end network."

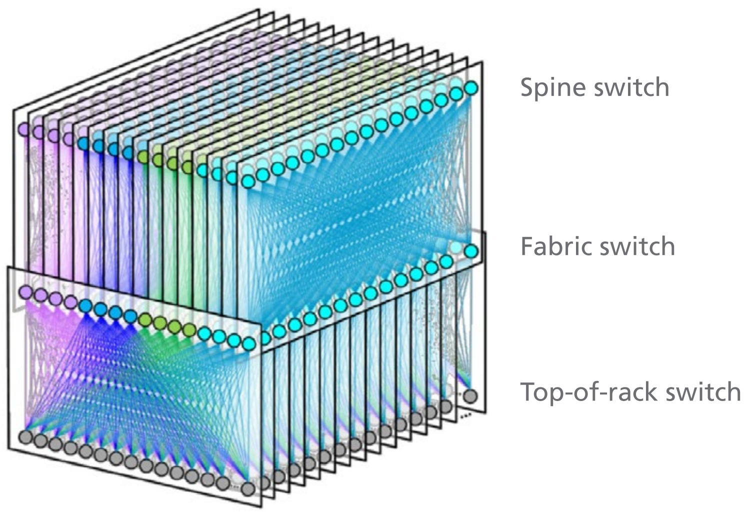

Fig. 2: FaceBook F16 data center network topology
(source: https://engineering.fb.com/2019/03/14/data-center-engineering/f16-minipack/)

Data halls with AI clusters

As noted, AI clusters have unique data processing requirements and thus require a new data center architecture. The GPU servers require much more connectivity between servers but, due to power and heat restraints, there are fewer servers per rack. Therefore, there is more inter-rack cabling in an AI data center than in traditional data centers. Each GPU server is connected to a switch within the row or room. These links require 100G to 400G at distances that cannot be supported by copper. In addition, each server requires connectivity to the switch fabric, storage, and out-of-band management.

Example: NVIDIA

As an example, we can look at the architecture proposed by NVIDIA, a leader in the AI space. NVIDIA's latest GPU server is the DGX H100 and has 4x800G ports to switches (operated as 8x400GE), 4x400GE ports to storage, and 1GE and 10GE ports for management. A DGX SuperPOD (Figure 3) can contain 32 of these GPU servers connected to 18 switches in a single row. Each row would then have 384x400GE fiber links for switch fabric and storage and 64 copper links for management. This is a significant increase in the number of fiber links in the data hall. The F16 architecture mentioned above would have 128 (8x16) duplex fiber cables with the same number of server racks.

What link lengths are in an AI cluster?

In the ideal scenario illustrated by NVIDIA, all the GPU servers in an AI cluster will be close together. AI/machine learning algorithms, like high-performance computing (HPC), are extremely sensitive to latency. One estimate claims that 30 percent of the time to run a large training model is spent on network latency and 70 percent is spent on compute time. Since training a large model can cost up to US$10 million, this networking time represents a significant cost. Even a latency saving of 50 nanoseconds, or 10 m of fiber, is significant. Nearly all the links in AI clusters are limited to 100 m reaches.

Unfortunately, not all data centers will be able to locate the GPU server racks in the same row. These racks require around 40 kW to power the GPU servers. This is more power than typical server racks, and data centers built with lower power requirements will need to space out their GPU racks.

Fig. 3: Rendering of NVIDIA SuperPOD
(source: https://www.nvidia.com/en-us/data-center/dgx-superpod/)

Which transceivers should you use?

Operators should carefully consider which optical transceivers and fiber cables they will use in their AI clusters to minimize cost and power consumption. As explained above, the longest links within an AI cluster will be limited to 100 m. Due to the short reach, the optics cost will be dominated by the transceiver. Transceivers that use parallel fiber will have an advantage: They do not require the optical multiplexers and demultiplexers used for wavelength division multiplexing (WDM). This results in both lower cost and lower power for transceivers with parallel fiber. The transceiver cost savings more than offset the small increase in cost for a multifiber cable instead of a duplex fiber cable. For example, using 400G-DR4 transceivers with eight-fiber cables is more cost-effective than 400G-FR4 transceivers with duplex fiber cable.

Links up to 100 m are supported by singlemode fiber and multimode fiber applications. Advances like silicon photonics have reduced the cost of singlemode transceivers—bringing them closer to the cost of equivalent multimode transceivers. Our market research indicates that, for high-speed transceivers (400G+), the cost of a singlemode transceiver is double the cost of an equivalent multimode transceiver. While multimode fiber has a slightly higher cost than singlemode fiber, the difference in cable cost between multimode and singlemode is smaller since multifiber cable costs are dominated by MPO connectors.

In addition, high-speed multimode transceivers use one to two watts less power than their singlemode counterparts. With 768 transceivers in a single AI cluster (128 memory links + 256 switch links X2), using multimode fiber will save up to 1.5 kW. This may seem small compared to the 10 kW that each DGX H100 consumes, but, for AI clusters, any opportunity to save power will be welcome.

In 2022, the IEEE Short Reach Fiber Task Force completed work on IEEE 802.3db, which standardized a new multimode very short reach (VR) transceiver. The new standard targets in-row cabling like AI clusters with maximum reach of 50 m. These transceivers have the potential to offer the lowest cost and power consumption for AI connectivity.

Transceivers vs. AOCs

Many AI, ML and HPC clusters use active optical cables (AOCs) to interconnect GPUs and switches. An AOC is a fiber cable with integrated optical transmitters and receivers on either end. Most are used for short reaches and are typically paired with multimode fiber and VCSELs. High-speed (>40G) active optical cables will use the same OM3 or OM4 fiber as cables that connect optical transceivers. The transmitters and receivers in an AOC may be the same as in analogous transceivers but are the castoffs. Neither the transmitter nor receiver must meet rigorous interoperability specs; they only need to operate with the specific unit attached to the other end of the cable. Since no optical connectors are accessible to the installer, the skills required to clean and inspect fiber connectors are not needed.

The downside of AOCs is that they do not have the flexibility offered by transceivers. Installing AOCs is time-consuming, as the cable must be routed with the transceiver attached. Correctly installing AOCs with breakouts is especially challenging. The failure rate for AOCs is double that of equivalent transceivers. When an AOC fails, a new AOC must be routed through the network. This takes away from the compute time. Finally, when it is time to upgrade the network links, the AOCs must be removed and replaced with new AOCs. With transceivers, the fiber cabling is part of the infrastructure and may remain in place for several generations of data rates.

Conclusion

Careful consideration of the AI cluster cabling will help save cost, power, and installation time. The right fiber cabling will enable organizations to fully benefit from artificial intelligence.

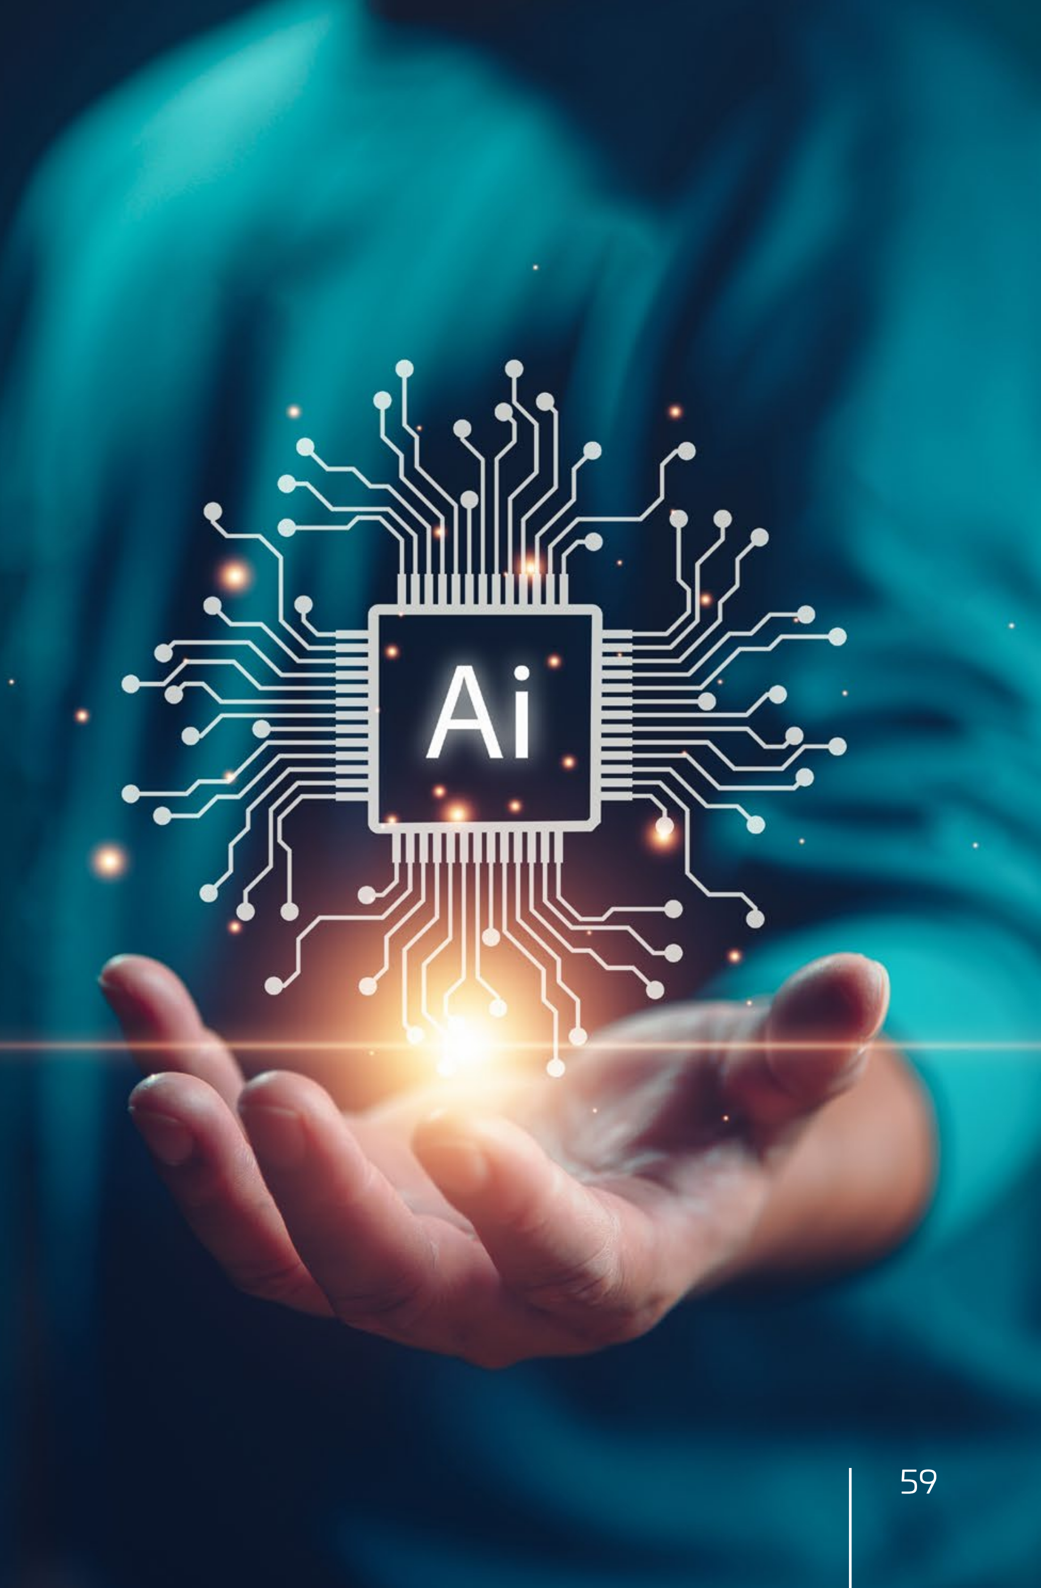

COMMSCOPE®

commscope.com

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

EB-117988-EN (11/23)