

Re-tooling the data center to accelerate deployment

The cost and complexity of designing and deploying modern data centers are growing. The trend is being driven by multiple factors that all share a common denominator: the need to transport, process and store the vast amounts data being generated by business, industrial, networking and social applications. The technologies being developed to accomplish these tasks are creating daunting challenges for data center managers, especially at the hyperscale and cloud-based level.

Increasing use of machine-to-machine computing is driving the accelerated adoption of more fiber. Higher fiber core counts are being pushed deeper into the network and are no longer necessarily the reserve of the core network.

As seen in Figure 1, a new generation of network switches and server platforms is providing higher density compute and transport capability per rack unit. Driven by improved chip technology, switch performance has increased dramatically. In 2014, switch bandwidth maxed out at 3.2 Tbps; three years later, 16 nm silicon gate technology doubled the maximum bandwidth to 6.4 Tbps. Today's switched chip technologies deliver 12.8 Tbps of bandwidth, enabling the use of 256 50G ports, 128 100GbE ports or 32 400GbE ports.



Figure 1: Chip I/O and optical integration
Source: Microsoft Cloud Hardware Futures—OCP 2019

At the end of 2019, Broadcom, the leading merchant silicon vendor, released its latest version of switching technology, Tomahawk® 4 (TH4). TH4 increases the switch chip bandwidth from 12.8 Tbps to 25.6 Tbps with proportional increases in the

chip radix as well. As a result, it delivers more bandwidth to the faceplate of the switch chassis.

To keep pace with the increase in chip technology, data centers are having to re-think their transport media. With standard twisted-pair copper having peaked at 10GbE, optical fiber is the only viable option to support the new higher speeds. However, as Figure 2 indicates, even optical technologies are struggling to keep up with faster switched chips.

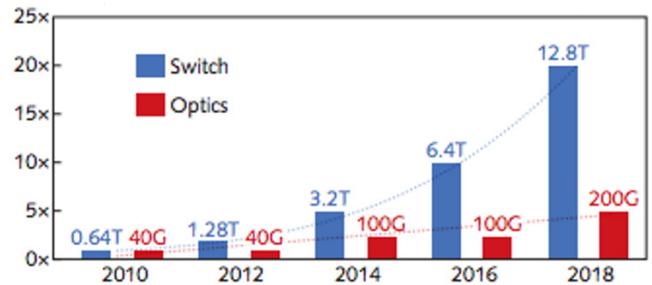


Figure 2: Relative bandwidth increase since 2010 (= 1x)
Source: Kucharewski et al, OFC M3J.7 2018

The optical industry is responding with new technologies and standards. VCSELs (vertical cavity surface emitting lasers) are now able to modulate at above 25 Gbps. Adding PAM4 (pulse amplitude modulation-4 level) encoding enables VCSELs to increase optical channel transmission to 56 Gbps over multimode fiber. 50 Gbps enables 400G transmission, which is seen as the next frontier in the evolution of hyperscale data center speeds.

Network complexity threatens ROI

Figure 2 illustrates how quickly relative bandwidth in the data center is accelerating. Against this background of increased

complexity, data center managers and owners are under pressure to increase optical density while keeping their fiber cabling manageable and adaptable to future changes. The demand for increased fiber deployment creates a cascade of challenges, with one of the most important being the time and resources required to bring new data center space online.

According to Enterprise AI, “Hyperscale has fundamentally transformed the data center market. Its scale and complexity have had a similarly profound impact on what we call the ‘value chain,’ the ecosystem of people involved in bringing such immense capacity to life.”

For data center managers and installation contractors, compressed timeframes and a shortage of skilled personnel are combining to create a perfect storm, with costly implications. CapEx spending from the top hyperscale operators totaled a record \$120 billion in 2018.ⁱⁱⁱ According to data from the U.S. Chamber of Commerce, infrastructure represents 82 to 85 percent of the CapEx spend, with labor and installation accounting for 75 percent of all infrastructure costs.^{iv}

“As the amount of data, distributed applications and other real-time capabilities emerge, infrastructure vendors also reported growing pressure to complete projects on tighter schedules.”ⁱⁱ

The ability to reduce infrastructure deployment times without sacrificing network design, performance or reliability can go a long way toward tilting the financial equation in a more positive direction. With labor and installation commanding the lion’s share of CapEx costs, accelerating the build-out of the physical layer—from design to commissioning—helps reduce total cost of ownership. But reducing deployment time isn’t just about cutting costs; the faster data center managers and owners can bring new capacity online, the stronger the bottom-line revenue.

The good news is that there are many opportunities during the planning and installation of the cabling infrastructure that can help reduce time to market and boost ROI. This paper examines some of the key challenges and opportunities involved.

Ordering and installation

Today’s data center owner or operator typically has a 12-week window in which to design and deploy the required physical network infrastructure. With millions of dollars of potential revenue at stake, the infrastructure must be turned up on time—regardless of the project’s complexity. However, the precise start

date for that short window is fluid. Compounding the challenge is that many sub-components used in the build have projected delivery dates that can fall on either side of that window.

During the deployment process, skilled installation teams trained in optical splicing and termination can make up some lost time; but, because they are in high demand, they are not always available. As Figure 3 indicates, the shortage of skilled labor is prevalent across the entire value chain and is seen as one of the four biggest challenges for data centers and their contractors.ⁱⁱ Assuming installation teams are available, the difficulty in aligning their busy schedules with the delivery of the infrastructure components further reduces the margin of error during the 12-week window.

Biggest challenges with hyperscale and colocation data center projects

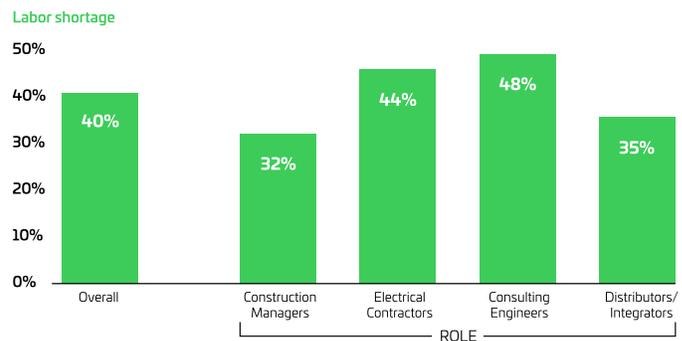


Figure 3: Labor shortages across the data center value chain

Source: Schneider Electric

Even if the fiber and installers arrive at the same time, the chances of completing the installation are shrinking due to the increasing number of fibers required. Fifteen years ago, most fiber backbones in the data center used no more than 96 strands, including coverage for diverse and redundant routing. Current fiber counts of 144, 288, and 864 are becoming the norm, while interconnect cables and those used across hyper- and cloud-scale data centers are approaching 3,456 strands. Terminating the higher fiber count cables requires more time. Cleaving and splicing a 144-core count cable with 12 MPO pigtailed can take 2.5 hours per end. To terminate both ends of one 144-core cable requires more than half a day. The time required to complete the cable termination increases as the optical core count of the fiber cable grows.

One other circumstance that makes any traditional data center deployment more complex is that these environments are usually high-security buildings. Data center managers are highly protective of their “vaults.” Everyone entering the data center rooms must pass through a thorough security scan. If you need a tool from the truck, need to go to the restroom, or need a cup of coffee, you must go through security.

Traditional installation methods require large crews, so even the routine process of entering the data center can become cumbersome—eating up valuable time. Fortunately, there are opportunities for installation crews to make up some of that lost time.

Advantages of pre-terminated fiber cabling

Pre-terminated, factory-connected cables significantly reduce the time to install and connect fibers to switches and routers. Fibers are spliced and terminated in a controlled environment. All end faces are inspected, and the insertion losses are verified, before the cabling leaves the factory. When the fiber arrives on site, installers can immediately start rolling it out into the raceways, using the factory-measured insertion loss as a reference when testing the channels.

Using pre-terminated versus field-terminated cabling enables installers to reduce deployment time by five hours for every 144-fiber cable installed.

Using pre-terminated connectivity—like MPO24, MPO16, MPO12 and MPO8—ensures first-time installation of optical channels with known guaranteed optical loss budgets. It further assures that the optical channels deployed meet the performance levels specified during the design. By reducing or eliminating the need for optical loss testing and remediation, deployment teams can further reduce installation times.

Using pre-terminated versus field-terminated cabling enables installers to reduce deployment time by five hours for every 144-fiber cable installed. When multiplied by the total number of cables used in the data center, the time savings quickly add up—providing more time for the integration and turn-up phases.

Day 2 cabling management

As the fiber port counts increase from a few hundred to several thousand, the sprawling cabling network has become significantly more difficult to manage. After cabling installation and turn-up, higher fiber counts often continue to negatively affect costs and operations. Moves, adds and changes are more challenging as technicians struggle to identify which fibers run between specific servers and switch ports. Should an issue arise causing service disruption, mean time to repair will be affected as well.

Other Day 2 operational challenges focus on the sheer speed of change within the data center. Throughput speeds are increasing,

standards are evolving, and more services and networks are being interconnected or converged. To keep pace, managers may “rip and replace” or layer new networks on top of existing networks—all while trying to meet their service level agreements (SLAs).

Without a cohesive connectivity management platform, the challenge of managing the inside cabling plant rises exponentially as more fiber strands are deployed. One of the biggest areas of concern is how to manage the growing amount of fiber patch cord slack. There are several solutions to help data centers manage, route and organize their sprawling cabling plant.

OPTICAL DISTRIBUTION FRAMES (ODF)

To address the issue of excess fiber patch cord material, data centers are turning to the use of optical distribution frames (ODFs) for cross-connect and interconnect applications. ODFs provide high-density support and improved cable management. A logical management routing system keeps all the fibers identifiable and accessible for more rapid installations, turn-ups and maintenance. Unlike standard rack panels, ODFs are modular and agile—specifically designed for fast-changing, fiber-dense environments. The solutions consist of the ODF frame and a variety of modular chassis, adapter packs and modules.

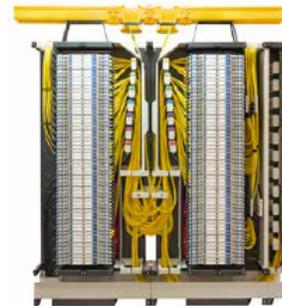


Figure 4: Two fully-loaded ODFs, side by side

Today's ODF solutions provide quick access to cross-connect port fibers and easy-to-follow cable routing paths—enabling technicians to locate and trace individual fibers between different halls and racks. The result is accelerated installation of new fibers and services, and improved handling of moves, adds and changes.

HIGH-DENSITY AND ULTRA HIGH-DENSITY PANELS

As connectivity moves from the core to the middle of the row or the server cabinets, focus is shifting to the ability to quickly and efficiently patch these circuits. Choosing patch panel solutions that are scalable, modular and flexible is important; having the ability to support increasing fiber counts while keeping the fibers secure and accessible is, perhaps, even more important.

High-density (HD), ultra high-density (UD), and enhanced high-density (EHD) fiber patch and splice panels enable data centers to support the growing number of in-row and server cabinet connections characteristic of today's leaf-and-spine networks. Some EHD panels provide up to 72 duplex LC or MPO ports per RU—singlemode or multimode. A key benefit is the ability to migrate from 10G serial to 100/400G parallel transmission without expanding equipment footprint.



Figure 5: HD and UD patch panels from CommScope

The UD, HD and EHD panels are also designed for faster installation and reduced MTTR. Tool-less install of rear trunk cables significantly reduces deployment time. An innovative sliding split-tray design provides open access to all fibers while protecting active links during modifications. A full line of fiber modules and adapter packs supports multiple shelf platforms for better agility and scalability.

If the average length of a cross-connect patch cord is three meters (roughly 9 feet), a 1,200-fiber port cross-connect (modest in today's larger data centers) using LC patch cords will require 600 patch cords and 1,800 meters (nearly 1 mile) of patch cord material to be dressed and managed.

FIBER ROUTING

Fiber raceway systems route fiber-optic patch cords and multi-fiber cable assemblies to and from fiber splice enclosures, fiber distribution frames and fiber-optic terminal devices. More than a convenience for keeping fiber cabling out of the way, raceways play a critical role in maximizing optical performance by ensuring fibers do not exceed their maximum bend radius.

Most fiber routing systems consist of modular components—horizontal and vertical runs and elbows, downspouts, junctions, and flexible exit kits—to enable a variety of deployment options. Some also include software-based design tools to ensure the

raceway conforms to load and bend radius restrictions. Fast, tool-less installation is an important (but often overlooked) benefit, especially given the short 12-week window for ordering and deploying the cabling infrastructure.

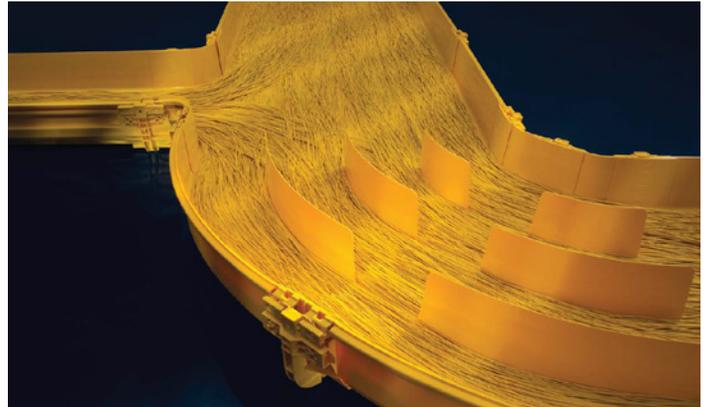


Figure 6: FiberGuide® fiber raceway system

Pre-provisioning services

In addition to the physical cabling and connectivity solutions, many infrastructure providers offer a variety of make-ready services to help data center managers further reduce deployment time as well as time spent to manage the infrastructure.

RACK AND STACK

Services like rack and stack involve pre-provisioning of the data center racks off-site. All the necessary networking components are assembled and tested ahead of time. When the build team is ready, the racks can be rolled into position and connected to the cabling infrastructure, and then the backbone cabling. Using this approach enables the data center to turn up a network in hours rather than days or weeks.

FIBER CHARACTERIZATION

Fiber characterization is a process of validating that an installed fiber connection will support a given set of applications before “lighting it up” with production data. It’s a critical step—not only when expanding or upgrading the network but during the initial build-out as well. The process of fiber characterization can involve multiple steps:

- **Validation:** Ensure the delivered product (fiber and connector types) not only matches the invoice but is in good physical condition. This step also involves checking to make sure the fiber has been properly prepared, and the end-faces are clean and properly cut and spliced.

- **Testing:** Confirm the fiber meets critical performance attributes. These tests include optical time-domain reflection (OTDR), chromatic dispersion (CD), polarization mode dispersion (PMD), insertion loss and optical return loss.
- **Documentation:** Create a complete record of the cabling infrastructure. This includes route and span drawings, cable and duct sizes, key splice points, utilization by service or customer, and end-to-end dynamic traces/light paths. Thorough documentation will ensure easier and faster Day 2 servicing and troubleshooting.

For larger deployments, fiber characterization may require third-party management to help coordinate the efforts of the data center’s WAN circuit acquisition teams along with the technical personnel from the fiber service provider.

Tools for expediting link design

While the actual process of designing the end-to-end links occurs prior to the components being ordered and installed, the accuracy of the design can significantly impact deployment times. Today’s complex modulation speeds and faster link speeds are more sensitive to interference. Shrinking loss budgets place increasing demands on designers to estimate link performance with greater accuracy.

For those installing fiber-optic systems, it is critical to have a reference of the optical losses they should expect. For the

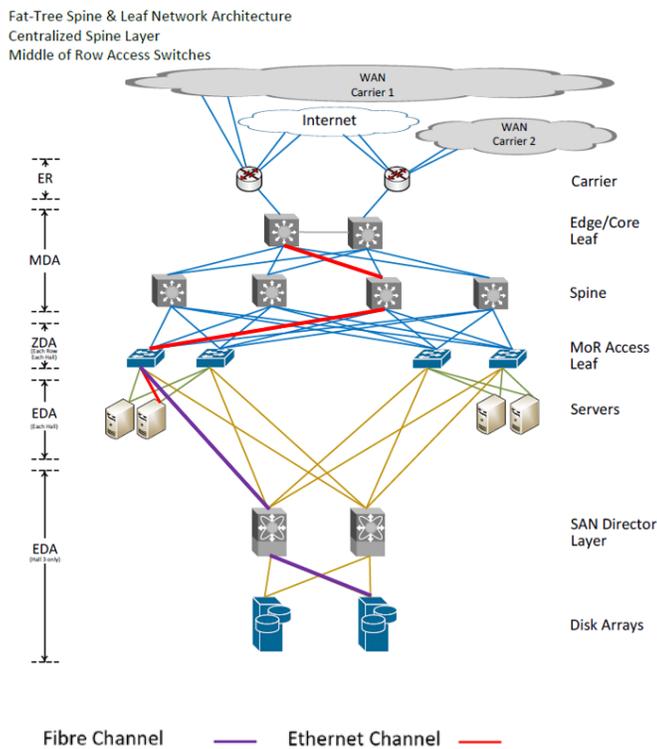


Figure 7: The IT department logical design

cabling to comply with customers’ expectation and standard requirements, the installer must check and see if the measured performance—insertion loss specifically—is below the theoretical losses. To accomplish that, you need a thorough understanding of how the losses of different components (e.g., cable, jumpers and connectors) intertwine to manually calculate the total link losses.

Various link design tools enable engineers to simulate the optical performance of end-to-end links in order to understand what applications specific designs will support. Using these tools, engineers can model a variety of link designs (see Figure 7) including multimode and singlemode transmission, low-loss and ultra low-loss cabling, connectors and components over a variety of distances.

Link design solutions enable a variety of specific functionalities for troubleshooting (including reflectometry in the time domain) to display a very clear map of the NEXT and RL behavior along the link. Additional functionalities include sophisticated interchangeable copper and fiber-optic modules, full-color touch screen, and detailed graphics of each parameter to be measured.

Full-featured design tools help ensure the signal path will perform as needed, before the first component or cable is ordered. This not only saves time and money during installation, it helps eliminate any surprises at the most critical time—during turn-up.

Reducing fiber counts while increasing capacity

As part of the link design process, engineers now also have more options for delivering increased capacity without adding more fiber. The list of modulation techniques for enabling wavelength multiplexing continues to grow. OM5 wideband multimode fiber (WBMMF), for example, expands the specified high-bandwidth operating wavelength spectrum, from a single wavelength at 850 nm to a range of more than 100 nm, including 850 nm through 950 nm. This wide specification range enhances the capability of shortwave division multiplexing (SWDM) to transmit 40G and 100G over a single pair of multimode fibers. The result, as illustrated in Figure 8, is a dramatic increase in the capacity of

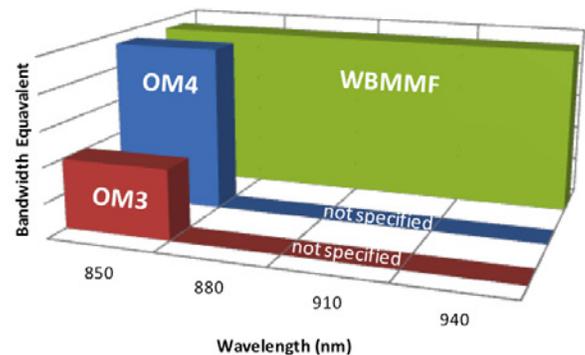


Figure 8: Bandwidth comparison of OM3, OM4 and OM5 WBMMF

parallel-fiber infrastructure—opening the door to four-pair 400GE and future 800GE applications as standards emerge.

Another modulation technology quickly gaining favor is four-level pulse amplitude modulation (PAM4), which is recommended by the 400GbE standards as the preferred modulation format for serial 400GE data center interfaces. Compared to the two-state non-return-to-zero (NRZ) modulation used in 100GE, PAM4 (see Figure 9) effectively doubles the data rate for a link bandwidth at the expense of reduced signal-to-noise ratio (SNR).




Figure 9: PAM4 modulation versus NRZ
Source: Mellanox Inc., 2017

Putting the expanding toolbox of solutions to work

As latency demands become more stringent and data centers adopt the mesh connectivity of spine-and-leaf networks, optical fiber counts are increasing exponentially. This is putting more pressure on data center operators and owners to reduce the time and cost needed to deploy, maintain and manage the growing fiber plant without sacrificing scalability and optical performance.

The good news is that the industry is responding with better-designed solutions—enabling installation and technical maintenance teams to work more efficiently. From higher density ODFs and patch panels to higher bandwidth fiber, easier cable routing and a variety of pre-provisioning services, data center managers have more tools at their disposal.

References

- i Broadcom Ships Tomahawk 4, Industry's Highest Bandwidth Ethernet Switch Chip at 25.6 Terabits per Second; Broadcom, press release; December 9, 2019.
- ii Hot Demand for Hyperscale Data Centers Hampered by Skills Gap; Enterprise AI magazine; October 17, 2019.
- iii The Impact of Hyperscale Data Centers: How the Wave is Changing the Value Chain; Schneider Electric, research report; October 2019.
- iv Data Centers: Jobs and Opportunities in Communities Nationwide; U.S. Chamber of Commerce Technology Engagement Center, industry report; 2017.

COMMSCOPE®

commscope.com

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

© 2020 CommScope, Inc. All rights reserved.

Unless otherwise noted, all trademarks identified by ® or ™ are registered trademarks, respectively, of CommScope, Inc. This document is for planning purposes only and is not intended to modify or supplement any specifications or warranties relating to CommScope products or services. CommScope is committed to the highest standards of business integrity and environmental sustainability with a number of CommScope's facilities across the globe certified in accordance with international standards, including ISO 9001, TL 9000, and ISO 14001. Further information regarding CommScope's commitment can be found at www.commscope.com/About-Us/Corporate-Responsibility-and-Sustainability.